Universiteit Leiden
The Netherlands

ICT in Business and the Public Sector

Disruptive technology and firm performance:
Analyzing disruptiveness in firm patent portfolios

Name: Ruben Klijn
Student-no: s1681907

Date: August 11, 2020

First Supervisor: Jian Wang
Second Supervisor: Frank Takes

Master's Thesis

Leiden Institute of Advanced Computer Science (LIACS)
Leiden University
Niels Bohrweg 1
2333 CA Leiden
The Netherlands

Abstract

Innovation plays an important role in company performance. A commonly asked question is whether a company should exploit the knowledge it already has or explore in search of something completely new. Previous works that study the balance between conducting exploring and exploiting innovations assume a relatively simple dichotomy between exploring and exploiting innovations. This research tries to fill the gap in the literature by making use of a sophisticated method (incorporating a patent citation network) of determining how explorative (disruptive) or exploitative (stabilizing) innovations really are on a continuous scale. This allows for looking at the true composition of innovations within firms, something that has not been done before in this context. Furthermore, this research uses a large scale panel data set constructed from multiple data sources to find the distribution of exploring and exploiting innovations that is optimal for firm performance. In order to achieve this, 21 years worth of patent data from six different technology fields are utilized as a way of capturing and analyzing innovation. By making use of a large patent citation network (consisting of over 1.9 million patents (nodes) and more than 7.6 million edges), each patent is classified on a continuous scale from -1 to 1. Here, -1 means that the patent is maximally stabilizing and 1 means that it is maximally disruptive. This results in an unbalanced panel data set including 415 companies and 3414 observations. By using fixed effects OLS and quantile regression models to analyze firm patent portfolios, this research concludes that having a relatively large number of stabilizing patents while, at the same time, having only a few destabilizing patents increases innovative performance. It is also observed that investments in disruptive innovation make sure that the company hits a good baseline in terms of minimal financial performance, while not significantly raising the chance of reaching extreme success. Furthermore, a patent portfolio that is more varied in terms of disruptiveness is likely to result in higher innovative performance.

Table of Contents

1 Introduction

Innovation is important for good reasons: most industry leading companies tend to be innovative and embrace innovation to large extends [Purcell, 2019]. The importance of innovation to firms has been studied extensively on different levels like firm productivity [Rao et al., 2001] and overall firm performance [Wang and Wang, 2012]. Nevertheless, it is not enough to just state the importance of innovation as companies also need to know how they should innovate, which is the focus of this work.

In order to research this, the extend to which a company is innovative needs to be measured. A good and established way of measuring innovation is by analyzing patents [Abraham and Moitra, 2001]. Besides, innovation can generally be classified as either exploring or exploiting. Exploration of ideas and concepts is based on elements such as discovery, experimentation and innovation, while exploitation of ideas and concepts is based on efficiency, production and execution [March, 1991]. Exploring patents (also known as disruptive patents) are about novelty, doing something that has not been done before, and exploiting patents (also known as stabilizing patents) are about exploiting prior knowledge and experience.

Given these two types of patents, what should a company do? Should it invest in more exploiting inventions or exploring inventions in order to optimize firm performance? This question is addressed in some scientific works ([Artz et al., 2010] [Uotila et al., 2009] [March, 1991] [Gupta et al., 2006]). In general, these works conclude that it is vital for a company to invest in both exploitation and exploration. Within existing literature, patents can be looked at in one of two ways:

1. A patent can either be fully exploring or fully exploiting. In this case, research can be performed on the share of exploring patents a company owns ([Verhoeven et al., 2016] [Geerts et al., 2017]).
2. A patent can be exploring or exploiting to a certain degree. Here, research can be performed on the the distribution of the degree [Funk and Owen-Smith, 2016].

The former method has been used in order to research the amount of exploring and exploiting activities that is optimal for firm performance [Belderbos et al., 2010]. However, it assumes a simple dichotomous way of patent novelty classification. This work offers an improvement by using a more sophisticated way of determining patent novelty on a continuous scale by incorporating a patent citation network [Funk and Owen-Smith, 2016]. This allows for studying the composition of patents within a patent portfolio in detail to look for the characteristics of patent portfolios that result in optimal firm performance, which is something that previous works have not achieved yet. To study this, the following research question is made:

What is the optimal distribution of exploring and exploiting patents that will maximize a company's innovative and financial performance?

To answer this research question, concepts and techniques of several fields are combined. Firstly, the field of network science plays an important role due to incorporation of a patent citation network. Secondly, there are databases available that hold large amounts of patent data and firm performance data. So, knowledge from the field of data science is also applied in order to construct and work with a large data set from various databases. Thirdly, the field of economics/business is vital to extract knowledge from the data and to answer the research question. Therefore, this work combines techniques from network science, data science and economics/business in order to answer the research question, which helps this work to stand out from other works.

By making use of 21 years of patent data, patent portfolios of a large number of companies are constructed. For each of these patents, their disruptiveness on a continuous scale is calculated using a large patent citation network consisting of over 1.9 million patents (nodes) and more than 7.6 million edges. This results in a data set consisting of 415 companies and 3414 observations in total to experiment with. By using fixed effects OLS and quantile regression models, patent portfolios are analyzed to find characteristics that result in optimal innovative and financial firm performance. Financial performance is measured by Return On Assets (ROA), innovative performance is measured by looking at citations and the number of patents in a patent portfolio.

The work concludes that, in general, owning more exploiting patents while having a relatively small number of exploring patents is likely to result in better firm performance. Also, a more varied patent portfolio in terms of disruptiveness leads to improved innovative performance, while having relatively more disruptive patents results in lower innovative performance. Additionally, companies in the lower quantiles benefit most from having more exploring innovations as it allows them to obtain a good minimum of financial performance. However, this does not significantly improve the chances to achieve extreme success.

This work adds to existing theory by providing new knowledge by applying a sophisticated patent classification measure combined with a relatively large data set. In terms of business practice, this works provides added insights on what the optimal patent portfolio should look like regarding various firm performance indicators.

The document starts with a literature review after which methods and data used in performing the research are explained in detail. Then, the experiments and results are discussed before ending with a conclusion.

2 Literature Review

This chapter consists of a literature review to shed some light on various related topics. First, general works that utilize patent analysis are mentioned, after which literature on the exploitation-exploration problem related to firm performance is discussed. This is followed by a description of various patent classification measures and a brief statement about the positioning of this work.

2.1 Patent Analysis

Patents serve the purpose of protecting intellectual property. In modern day and age, however, patents get used in different ways. For example, they serve the purpose of being a tool for guiding companies to improved performance and competitive advantages [Rivettte and Kline, 2000]. This shows that patent data can be used and studied in various ways. Some of the ways patent data can be looked at and used in research are mentioned next.

2.1.1 Bibliographic And Value Creation Approach

There exist two approaches for researching patent data: a bibliographic approach and a value creation approach [Grimaldi et al., 2015]. Patents are a source of bibliographic patent information that consists of elements like the entity that owns the patent, the people that created the invention, citations and patent classifications. This allows for the construction of, for example, patent citation networks that can be analyzed and improved.

An example of the application of the bibliographic approach consists of bridging the gap between science and technology. This can be done by the analysis of scientific references that occur within patents in order to find patterns between science creation and technological development [Verbeek et al., 2004]. Furthermore, the bibliographic approach can also be extended with other techniques from different fields such as text mining. This was applied to create a new kind of document retrieval system with improved search accuracy in order to improve patent analysis [Liu et al., 2011].

These examples of the application of the bibliographic approach in scientific research were given to emphasize one of its disadvantages: it is difficult to use the bibliographic approach on its own to enhance company performance directly. If, however, the bibliographic approach is extended with information and techniques from elsewhere (for example, with the help of computer science techniques as was the case with the last example), it becomes easier to apply it directly to a business context.

The value creation approach revolves around utilizing the strategic and economic value that patents can hold. Often, the end goal is to improve a company's value by studying the effects that technology and innovation have on company performance. This can be achieved by combining bibliographic patent data with data from other databases. It can include the addition of financial data, but also company descriptive data or forms of historical data.

This research makes use of both value creation approach and the bibliographic approach. The bibliographic approach is applied by studying the disruptiveness of patents using a patent citation network, while the value creation approach is applied in analyzing the effects of patent portfolios on company performance.

2.1.2 Network Science

The field of network science is involved in the analysis of patent data in numerous other works. The nature/composition of bibliographic patent data allows for the creation of several types of networks. One of these types consists of co-authorship networks, which allows for studying relations between the co-authors/inventors of patented innovations. Such a network gives possibilities for the analysis of knowledge flows [Breschi and Lissoni, 2005], the collaboration process involved in creating patents [Bergek and Bruzelius, 2010] or firm performance related to firm collaboration [Powell et al., 1999].

Another type of network involved in patent analysis is the citation network in which links are formed between patents that cite (or get cited by) other patents. There are numerous works that analyze patents this way in order to, for example, study the diffusion of technologies [Chang et al., 2009], the prediction of emerging technologies [Erdi et al., 2012] and technology evolution over time [Martinelli, 2010]

2.1.3 Data Science

Besides constructing networks, there are other ways of extracting and analyzing patent data. For example, text mining approaches proved to be another valuable method of analyzing patents. Such methods have been used for monitoring emerging technologies [Joung and Kim, 2017], analyzing firm patent portfolios [Fattori et al., 2003] and patent evaluation [Han and Sohn, 2014].

No matter the approach of analysis of patent data, having sufficient data is important. It is not uncommon in patent analysis studies to incorporate relatively small data sets, often focused on a particular type of innovation (like LED technology [Choi and Hwang, 2014]), technology field or business sector (like the wellness-care industry [Kim and Bae, 2017]). However, especially when working with network data, the boundary specification problem should be taken into account when choosing a subset of data to make sure that results are reliable [Laumann et al., 1983] .

2.2 Exploration, Exploitation And Firm Performance

The link between exploration, exploitation and firm performance is explained next, together with other literature related to the subject. After that, light is shed on works about determining the balance between exploitation and exploration related to firm performance.

2.2.1 The Importance Of Balance

It is important to know why companies need to find a balance in exploration and exploitation. A company should devote sufficiently enough resources to exploration to achieve good future performance, while simultaneously, enough resources should go to exploitation to ensure good current performance [March, 1991]. This is explained by two phenomena: the success trap and the failure trap [Levinthal and March, 1993]. The success trap entails that a company focuses on exploitation too much because of the short-term successes that come along with exploitation. This behavior can cause obsolescence due to the fact that long-term profits (both in terms of money and knowledge) that come from exploratory activities are ignored.

On the contrary, the failure trap occurs when a company focuses too much on exploration. Exploratory activities have a high risk of failure. After such a failure, the company that focuses on more exploratory activities invests in more exploration because of the high-risk high-reward nature , which is likely to fail and initiate more exploration. This vicious circle repeats itself over and over again while no profits are being made, which can be detrimental to a company.

Other works have also concluded that a mixture of both exploring and exploiting research is needed to optimize technological performance [Geerts et al., 2017]. Besides, an imbalance in patents where there is too much focus on exploratory patents leads to bad company performance, as more money is spent on researching rather than making profit [Gupta et al., 2006], which is in line with the findings of Levinthal and March. So, it is important for a firm to become (simultaneously) ambidextrous, which means that it should participate in both exploring and exploiting activities at the same time, resulting in being more successful in the long run [Tushman and Iii, 2006].

According to many other works, there seems to be a positive relationship between ambidexterity and firm performance, while being either too much or too little ambidextrous as an organization is likely have negative consequences on firm performance [O'Reilly and Tushman, 2013]. Furthermore, there exist complex interaction effects between ambidexterity and many contextual factors like discipline, trust and support within organizations or the type of industry the organization operates in [Junni et al., 2013] [Gibson and Birkinshaw, 2004]. So, ambidexterity plays a big role in how a firm performs both in the short and long run.

2.2.2 Determining the balance

It has become clear that a good balance between exploitation and exploration within an organization is necessary. The next question then becomes what this balance should look like to improve firm performance. Some work has been done on this particular topic. For example, is has been observed that there exists a positive relationship between exploiting inventions and financial company performance, while there was a negative relationship between exploring inventions and financial company performance [Artz et al., 2010].

In addition, there exists a curvilinear relationship between a firm's degree of exploration and financial performance that is mediated by industry R&D intensity [Uotila et al., 2009]. Based on this work, research has been done that was partly devoted to finding the balance of exploring and exploiting technologies that is optimal for company performance. It concluded that having a 39% share of exploring patents (and, thus, having a 61% share of exploiting patents) results in the best performance for most companies [Belderbos et al., 2010]. Furthermore, during the process of classifying patents, they used a method that classifies a patent as either fully exploiting or fully exploring. They also used eight years worth of longitudinal data, incorporating data of 168 Japanese companies in their research.

2.3 Measuring Exploration And Exploitation

Patents can be used in various ways to measure the degree of exploratory and exploitative activities a firm participates in. One way of doing this is the "novelty in recombination" indicator, which makes use of International Patent Classification (IPC) classes [Verhoeven et al., 2016]. A patent can be classified with one or more IPC classes. According to this measure, if a patent has a combination of IPC classes that has been observed in earlier patents within the company over the last x years (for example, 10 years [Kovács, 2019]), a patent/innovation is familiar and, therefore, exploiting. If, on the other hand, that specific combination of IPC classes has not been observed before with other patents a company owns, the patent can be considered as novel (exploring).

A similar way classifying patents as exploiting or exploring is proposed by Geerts et al. and makes use of patent technology fields [Geerts et al., 2017]. According to this measure, if a company has not been active in a patent's technology field(s) in the previous five years, the patent is considered to be novel (exploratory). Additionally, if a company has only been recently in the patent's technology field (less than three years of experience), the patent is also exploratory due to a lack of experience in that field. Otherwise, the patent is classified as exploiting.

Lastly, there is also a more sophisticated measure of disruptiveness (novelty) called the CD_t *index* [Funk and Owen-Smith, 2016]. It calculates an index that determines to what extend a patent is consolidating or destabilizing, hence the name CD_t *index*. Specifically, it calculates a patent's disruptiveness on a continuous scale of -1 (fully stabilizing/exploiting/consolidating) to 1 (fully disruptive/exploring/destabilizing) based on a patent citation network. Conceptually, it works as follows: let's say there is a patent A. If future patents only cite patent A and none of the patents A itself cites, patent A is fully destabilizing since it renders the patents before it obsolete. If, on the other hand, future patents cite mostly the patents cited by patent A, patent A is classified as a stabilizing patent.

2.4 Positioning The Research

This research tries to fill the gap in the literature that exists in studying the effects of patent portfolios of companies and deviates from the already existing works in a couple of ways. First of all, this research makes use of a patent classification method that classifies patents using a continuous scale, meaning that a patent will be exploiting and exploring up to a certain extend. So, contrarily to other works, no binary patent classification method is used and no dichotomy is assumed in this research. Secondly, as said before, it is not unusual to work with a smaller subset of data in these kinds of works. However, a large data set is used in this research and is composed of data from multiple databases: 21 years of longitudinal patent data enriched with company financial data is incorporated. Lastly, current literature has not researched the optimal distribution of patents in detail yet within these limits. In fact, this work looks at the composition of patents within the patent portfolios of companies, which is something that has not been done before. This work will fill the gap in the literature by studying distributional statistics of patent portfolios that will maximize firm performance on a relatively large data set.

3 Data & Methods

The choices regarding data and methods that are utilized in this research are explained below. This chapter starts with an introduction to the patent novelty measure that is used in order to determine to what extend patents are stabilizing or disruptive. Then, the process of constructing the data set is shown in detail and dependent as well as independent variables are discussed. The chapter ends with an explanation of the method.

3.1 Patent Disruptiveness Measure

Fundamental to this research is the patent data on which the entire data set is based upon. Before telling in detail which patent data specifically is used, it is necessary to explain the method of assigning novelty scores to patents.

The novelty measure that was incorporated in this work is the one from Funk and Owen-Smith called the CD_t *index* [Funk and Owen-Smith, 2016]. This novelty measure is chosen because is provides an intuitive way of measuring a patent's degree of disruptiveness based on a patent citation network. In a citation network, forward citations can indicate importance of an innovation [Trajtenberg, 1990], whereas backward citations can measure how incremental the innovation is (in other words, how much it relies on previous knowledge) [Lanjouw and Schankerman, 2001].

Conceptually, the CD_t *index* is sound because it takes both forward and backward citations into account in calculating disruptiveness. It does this by determining to what extend a patent renders the patents that it cites obsolete based on the citations the patent receives relative to the patents that it cites.

In practice, the CD_t *index* classifies patents on a scale of -1 to 1, where a patent with a CD_t *index* of -1 means that the patent is maximally stabilizing, while a patent with a CD_t *index* of 1 is maximally disruptive. In this context, disruptive means that a patent is novel, something that has not been done before and challenges the status quo. Stabilizing means that a patent is more exploiting, utilizing what was already known before.

This measure makes use of a citation network $G = (V_1, V_2, V_3, E)$ This is a directed, acyclic graph. Let:

- E represent the edges in the graph.
- V_1 be the focal patent f, that is, the patent of which we want to calculate the CD_t *index*.
- V_2 be the set b of patents that are cited by the focal patent f
- V_3 be the set i of patents that cite f, any of the patents in b or both.
- t be the time.

Then, the CD_t *index* is calculated as follows:

$$CD_t = \frac{1}{n_t} \sum_{i=1}^{n} \frac{-2 * f_{it} * b_{it} + f_{it}}{w_{it}}, w_{it} > 0$$

where:

- $f_{it} = \begin{cases} 1 \text{ if } i \text{ cites focal patent } f, \\ 0 \text{ otherwise} \end{cases}$
- $b_{it} = \begin{cases} 1 \text{ if } i \text{ cites any patent of } b, \\ 0 \text{ otherwise} \end{cases}$
- n_t = the number of forward citations in i
- w_{it} is a matrix containing weights for patent i at time t denoting how important patent i is in the citation network

If a focal patent scores 1 according to this novelty measure, it means that the focal patent is maximally disruptive and, therefore, is focused on exploration. And if a focal patent scores -1, the focal patent is maximally stabilizing, which means that the focal patent is focused on exploitation. Lastly, if a focal patent scores 0, it is undefined whether the focal patent is focused on exploration or exploitation.

In this research, like Funk and Owen-Smith did in their work, the value of w_{it} will be held constant at 1 for simplicity. Determining appropriate values for w_{it} is out of scope for this research, but might be done in future work.

3.1.1 Example

Figure 1 shows an example of a patent citation network. The CD_t *index* for f is calculated as follows:

- For patent $i1$: $-2 * f_{1t} * b_{1t} + f_{1t} = -2 * 1 * 0 + 1 = 1$
- For patent $i2$: $-2 * f_{2t} * b_{2t} + f_{2t} = -2 * 1 * 0 + 1 = 1$
- For patent $i3$: $-2 * f_{3t} * b_{3t} + f_{3t} = -2 * 1 * 0 + 1 = 1$
- For patent $i4$: $-2 * f_{4t} * b_{4t} + f_{4t} = -2 * 1 * 0 + 1 = 1$

This leads to:

$$CD_t = \frac{1 + 1 + 1 + 1}{4} = 1$$

From this, we can conclude that the focal patent f is maximally disruptive and, therefore, is focused on exploration.

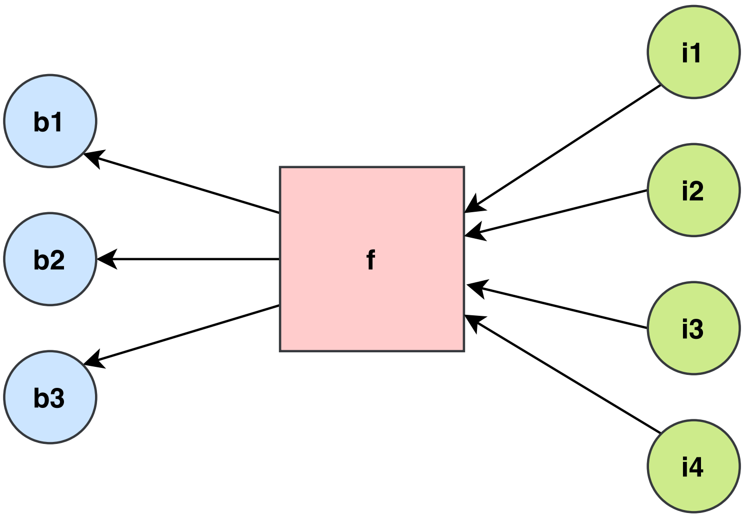

Figure 1: An example patent citation network were the focal patent f is maximally disruptive.

Figure 2 shows another example of a patent citation network. The corresponding CD_t *index* for f in this network is calculated as follows:

- For patent $i1$: $-2 * f_{1t} * b_{1t} + f_{1t} = -2 * 1 * 1 + 1 = -1$
- For patent $i2$: $-2 * f_{2t} * b_{2t} + f_{2t} = -2 * 1 * 1 + 1 = -1$
- For patent $i3$: $-2 * f_{3t} * b_{3t} + f_{3t} = -2 * 1 * 1 + 1 = -1$
- For patent $i4$: $-2 * f_{4t} * b_{4t} + f_{4t} = -2 * 1 * 1 + 1 = -1$

This leads to:

$$CD_t = \frac{-1 - 1 - 1 - 1}{4} = -1$$

From this, we can conclude that the focal patent f is maximally stabilizing and, therefore, is focused on exploitation.

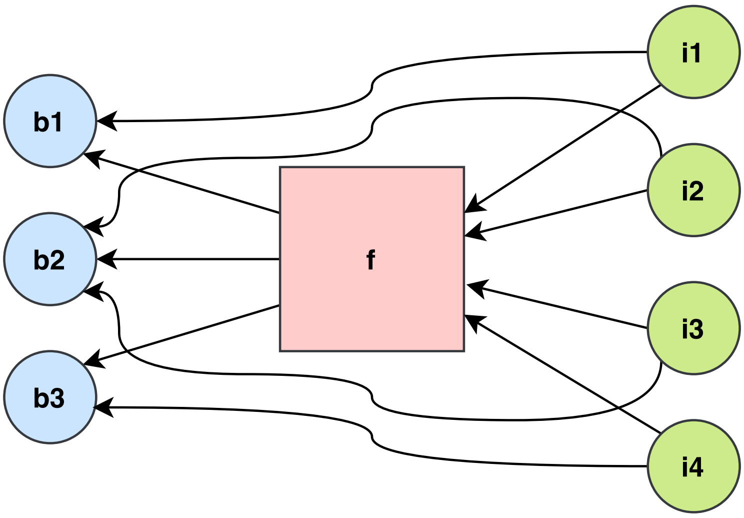

Figure 2: An example patent citation network were the focal patent f is maximally stabilizing.

3.2 Constructing The Data Set

The construction of the data set is divided into two general steps which are shown below. The first step is about the construction of the patent citation network and the second step sheds light on the way patents are linked to their companies.

3.2.1 Creating The Patent Citation Network

The disruptiveness measure explained earlier works with a patent citation network. In order to construct this network, numerous decisions regarding data selection need to be made. The raw patent data was retrieved from the PATSTAT database, maintained by the European Patent Office [EPO, 2020]. The PATSTAT database holds legal and bibliographical event patent data from countries worldwide, which makes it suited for (large scale) statistical analyses. From this database, a selection of focal patents, all patents that are cited by these focal patents and some other information about all patents involved (such as filing dates and technology fields) are retrieved as well. With this, it is possible to build a patent citation network.

For this research, the decision is made to use patents granted by the United States Patent and Trademark Office (USPTO) and filed from the beginning of 1980 until the end of 2000 as focal patents. In total, this yields 21 years of patent data, which results in a relatively large data set.

The patent data is further restricted by the use of technology field selection. A patent can be assigned to one or more technology fields to which it is most applicable. These technology fields are determined by the World Intellectual Property Organization

(WIPO) and the technology fields themselves used in this research originate from the IPC8 Technology Concordance table (updated last in July 2019).

Initially, this research focused only on focal patents assigned to technology field number 16, which is the field of pharmaceuticals. This is a fast moving, innovative field, which is why it is interesting for this research. The relationship between technology field 16 and other technology fields is shown in Table 1. This table shows the top-6 technology fields that are linked strongest to technology field 16 in terms of the number of times they got cited by the focal patents from technology field 16. The full table can be found in Appendix A.

For example, Table 1 shows that the focal pharmaceutical patents cite other pharmaceutical patents the most (650.173 times), while organic fine industry patents got cited by pharmaceutical patents second most with 400.342 citations. Please note that the number of links/citations to these fields is relatively high due to the fact that patents can be related to multiple technology fields.

Table 1: The Top-6 Technology Fields And Their Relationships

Technology Field	Field Name	Abbr.	Number Of Citations
16	Pharmaceuticals	Pharm	650.173
14	Organic Fine Industry	OFI	400.342
15	Biotechnology	Bio	278.733
13	Medical Technology	Med	119.625
11	Analysis Of Biological Materials	AOBM	68.882
19	Basic Materials Chemistry	BMC	66.609

The top-6 technology fields that showed the strongest relationship with technology field 16 (including their technology field names and abbreviations) based on the number of times patents from the technology fields got cited by the focal patents from technology field 16. Note that one patent can be assigned to multiple technology fields.

The decision is made to include patent data from technology field 11, 13, 14, 15 and 19 to the data from technology field 16 because of the fact that these fields appeared in the top-6 shown in Table 1. The strong relations between the fields make sense as they all include research related to medicine, chemistry, health care and technologies needed to create and use these inventions. Because these fields generate a large amount of patents, they are well suited for patent analysis purposes [Breitzman and Thomas, 2002]. This decision results in the patent citation network that contains:

- All patents filed at the USPTO from the beginning of 1980 until the end of 2000 that belonged to technology fields 11, 13, 14, 15, 16 or 19. These patents will be used as focal patents.
- For each focal patent, all of the patents cited by the focal patents are included without any restrictions.
- All patents from any other field that cite any of the focal patents in the data set are included without restrictions.
- Of course, the corresponding edges between patents are included to complete the citation network.

Using the patent citation network described above, the novelty measure of Funk and Owen-Smith is applied to all focal patents. However, some patents are older than others and, therefore, could have had more time to gather citations than patents that are newer. Moreover, research has shown that a peak in citations is reached within the first 5 years after introduction of a patent [Jaffe and Trajtenberg, 2002]. This is why the decision is made to include only relevant patents that are filed within a period of five years after the focal patent was filed in V_3 for the specific focal patent in V_1 in the calculations of the novelty scores. Like this, it allows for a fair way to calculate the CD_t *index* for each focal patent.

If we were to count all of the patents that are included in V_3 for the CD_t *index* calculations for all focal patents, the total accumulated number of patents in V_3 would be 113.710.830 without the five year restriction. With the five year restriction, however, this number is brought down to 113.661.036. This means that, with the five year restriction, $\frac{113661036}{113710830} * 100\% \approx 99.96\%$ of all possible patents in V_3 are used, which means that the five year restriction does not have a large impact on the results while at the same time confirming the findings from Jaffe and Trajtenberg.

3.2.2 Creating Patent Portfolios For Companies

The next step is to link the focal patents with valid novelty scores to the company to which they belong. For this, data provided by the National Bureau of Economic Research [NBER, 2020] was used. This data set provides the links between U.S. patent data from 1976 to 2006 and the companies to which they are assigned via Compustat. Now that the focal patents are linked to a company, it is possible to create a patent portfolio for each company: a patent portfolio is made up from the focal patents within the data set that belong to that company. Note that the patent portfolio for a company can vary each year. Now, the data set includes the patent portfolios for each company for all available years between 01-01-1980 and 31-12-2000 (which means that there is a panel data structure).

Lastly, for each year and for each company in the data set, company financial data is added for dependent variable construction. This financial data is retrieved from the Compustat financial database and includes financial data on the earnings before interest and the value of total assets for each company and each year [Compustat, 2020]. This results in an unbalanced panel data set, where each observation consists of a year worth of company data.

3.3 Dependent Variables

The four dependent variables are discussed below. First, three dependent variables for measuring a firm's innovative performance are explained before going over the dependent variable for measuring a firm's financial performance.

3.3.1 Innovative Performance

There exist numerous ways in which the innovative performance of a company could be measured. One way of measuring innovative performance is by counting the number of patents a company owns. Examples of works in which this measure is applied are the work of Furman and Hayes about studying innovative performance in economic growth areas [Furman and Hayes, 2004] and the work done by Beneito about the differences between in-house and contracted R&D performance [Beneito, 2006]. Other used ways of measuring innovative company performance include using the number of products that are patented [Laforet, 2008], the number of patents that are significant [Alsharkas, 2014], the market value of patents [Koski and Mäkinen, 2009] and R&D expenditures [Noori et al., 2017]. Each performance measure offers different perspectives of looking at innovative performance. Therefore, this work uses three different innovative performance measures that cover innovative productivity and quality.

Productivity

Firm productivity can be measured in various ways that incorporate different kinds of data. For example, methods exist that try to calculate productivity by taking firm financials and descriptive characteristics into account ([Dabla-Norris et al., 2012] [Chen et al., 2011] [Delgado et al., 2002]). However, a simpler way of measuring productive performance is counting patents, because patent counts are probably the most direct way of measuring firm productivity when it comes to innovation since it maps directly to the innovative output of a company [Pakes and Griliches, 1980]. Because of this, the patent count (number of patents) within a company's patent portfolio is going to be used as a dependent variable.

Quality

The average number of citations is also interesting to incorporate in this work due to the focus on quality of patents. There can be cases where a company has got ten patents in its portfolio, of which nine patents get cited only very few times while one patent gets cited a lot of times. In terms of total citations, the company has a decent innovative performance. However, when looking at average citations, it shows that the quality of all patents is not as good as the total citations make it seem. Therefore, to make a judgment on the relative innovative performance of a company, average citations of the patents within a company's patent portfolio is used as a dependent variable as well.

Besides this, the number of times patents owned by a company got cited is also used as a measure of innovative company performance. It has been shown in previous studies that counting patent citations is a valid way of measuring a company's innovative performance [Hagedoorn and Cloodt, 2003]. According to Hagedoorn and Cloodt, this depends on the basic assumption that there exists a positive relationship between the the number of citations a patent acquires and the quality of the patent. So, the total number of citations within a company's patent portfolio is also used as a dependent variable. However, bare in mind that this measure might be over-dominated by quantity of patents. To account for the company size effect, control variables for the number of patents and firm size in terms of the number of employees will be incorporated in this research.

Please note that for average citations, the number of patents and total citations, the log transformed variables are used. These variables are highly skewed and it is common practice to apply a log transformation in these cases. Upon further analysis, experiments showed that the R^2 value of the regression models increased when the log transformed variables were used compared to the standard, raw variables. The log transformation means that results should be interpreted differently: the models show changes to the dependent variable in percentages when the independent variable changes by 1.

3.3.2 Financial Performance

In other works, many measures for financial company performance are used. Some of these works make use of stock market valuation [Jayaraman et al., 2000], operating income ratios [Huson et al., 2004], Return On Assets (ROA) [Larcker et al., 2013] and Tobin's Q [Ghani and Ashraf, 2005]. Originally, the idea was to use Tobin's Q as a measure for firm performance. Although this firm performance measure is used frequently in other works, it is relatively hard to compute because it has a large number of components/variables that are needed in calculations. As a result, over half of the data set became unusable due to missing values for at least one of the components of the formula.

This meant that a simpler measure for firm performance was required. ROA is another firm performance measure that has been used extensively in other works and provides a way of reflecting the financial strength of a company in the long run [Baer and Frese, 2003]. ROA can be calculated in different ways and for this work, the method used by Fooladi and Chaleshtori was chosen because of its simplicity [Fooladi and Chaleshtori, 2011]. Within the context of this work, the following formula is used:

$$ROA = \frac{\text{Earnings Before Interest}}{\text{Total Assets}}$$

Unlike Tobin's Q, for all companies in the data a ROA value could be calculated. This is probably due to the fact that earnings before interest and total assets are values that can be found on almost every balance sheet, while Tobin's Q incorporates some variables that do not always need to be specified. Therefore, ROA was used as a dependent variable for measuring firm performance.

3.4 Independent Variables

In total, six independent variables are incorporated in the research, three of which are added as control variables. First, the three most important independent variables (which tell something about the patent portfolios) are explained, after which the three control variables are mentioned.

3.4.1 Patent Portfolio Distribution Statistics

For each company, for each year, basic information about patent portfolios is known. We know how many patents a company has within the patent citation network and we know of each of those patents the novelty score. This allows for building a distribution of novelty scores within the patent portfolio of a company. Consequently, information from these distributions can be captured and used to study their effects on the company performance. Specifically, the first, second and third moment of these disruptiveness distributions are used for this purpose.

The first moment gives the disruptiveness mean, the second moment gives the disruptiveness variance and the third moment gives the disruptiveness skewness. Remember that, if a company's patent portfolio distribution has a mean of -1, the company has patent portfolio that focuses on stabilizing patents, while a patent portfolio distribution with a mean of 1 shows that a company is more focused on disruptive patents.

Similarly, a low disruptiveness variance shows that the patents within the patent distribution have similar scores that are close to each other (this gives a more focused portfolio). A high disruptiveness variance indicates that the patents within the patent portfolio are more spread out and that the portfolio is less focused on a particular type of patent. Lastly, a disruptiveness skewness of -1 implies that the patent distribution is skewed left, while a disruptiveness skewness of 1 implies that the patent distribution is skewed right.

Distributional statistics are captured for each company for each year. However, it takes more than one year for a patent to have full effect on firm performance. Research has shown that there is approximately a two to three year gap between the "release" of a patent and its effect on firm performance [Ernst, 2001]. This shows that it is important to capture the effect of the patent portfolio distributions over time. Additionally, this research makes use of an unbalanced panel data set. Therefore, mean, variance and skewness of the portfolio distributions will be averaged over a three year time period.

For example, for company A in the year 1994, the disruptiveness mean independent variable is equal to the average disruptiveness values of patents in the portfolios at 1991, 1992 and 1993. The variance and skewness values are calculated in the same way, they are all statistics based on patents in the patent portfolios of the previous three years. If it is not possible for a company to calculate three year averages (for example, when a company occurred in the data set for only three years or less), the data is not considered during this research.

3.4.2 Control Variables

Three control variables are used in this research, namely the number of employees a company has, the number of patents a company owns and the average citations within a patent portfolio. The number of patents and average citations are lagged dependent variables. They are added as control control variables because they remain relatively stable through the years, meaning that the values of these variables in previous years are likely to affect the values in the current year. Therefore, a three year lag was introduced to these variables similar to the lagged disruptiveness mean, variance and skewness variables. Note that all of the control variables are log transformed.

3.4.3 Overview Of Variables

Table 2 presents an overview of the four dependent variables and the six independent variables and provides a short description.

Table 2: Overview Of All Dependent And Independent Variables

Variable	Description
Patents (ln)	The total number of patents a company has in the current year. (log transformed).
Avg. Citations (ln)	The average number of citations that come from the patents in a company's patent portfolio (log transformed).
Total Citations (ln)	The total number of citations that come from the patents in a company's patent portfolio (log transformed).
ROA	Return On Assets of a company.
Disr. Mean†	The disruptiveness mean (first moment) of a company's patent portfolio distribution of novelty scores.
Disr. Variance†	The disruptiveness variance (second moment) of a company's patent portfolio distribution of novelty scores.
Disr. Skewness†	The disruptiveness skewness (third moment) of a company's patent portfolio distribution of novelty scores.
Employees (ln)	The number of employees that a company has (log transformed).
Patents† (ln)	The total number of patents a company has (log transformed).
Avg. Citations† (ln)	The average number of times patents within the patent portfolio of a company got cited (log transformed).

An overview of both the dependent and independent variables with a short description. A † indicates that a variable represents statistics based on patents from the previous three years.

3.5 Experiments

Experiments in this work consist of both fixed effects OLS regressions and fixed effects quantile regressions for panel data. First, the fixed effects OLS regression experiments are explained before discussing the fixed effects quantile regression exepriments and restrictions to the data set.

3.5.1 Fixed Effects OLS Regression

A common way of studying the effects of predictors on company performance is to make use of OLS regression models. For example, such models are used for studying the effect of intellectual capital on firm performance [Clarke et al., 2011] and the effect of training of employees on firm performance [Hansson, 2003]. Within this research, an implementation of OLS regression suited for panel data is used and experiments are done with models that incorporate company fixed effects. Company fixed effects allow for researching the effect of patent portfolios on performance within companies, where elements like management style and company culture are relatively fixed. This work includes experiments that both include and exclude company (entity) fixed effects to compare within company differences as well as between company differences. Models are created for the four dependent variables of patents, average citations, total citations and ROA. This includes models created with and without entity fixed effects.

3.5.2 Fixed Effects Quantile Regression

Another way to study company performance is by using quantile regression models. To illustrate some of the uses of quantile regression, they have been used to study the effects of corporate social responsibility and corporate performance [Kang and Liu, 2014], but also innovative firm performance has been analyzed with quantile regressions [Ebersberger and Herstad, 2013]. Quantile regressions offer the advantage that the responses of independent variables can be studied at different quantiles of the distribution of the dependent variable. This means that it adds more depth to the experiments by, for example, comparing the effect of predictors on companies in the lower quantiles that perform worse with those in the higher quantiles that perform better.

The specific method used for doing quantile regression is an implementation of Koenker that is suited for panel data and incorporates fixed effects [Koenker, 2004]. Coad and Rao used a quantile regression approach to study innovation and firm growth in high-tech sectors [Coad and Rao, 2008]. In their approach, they used 10%, 25%, 50%, 75% and 90% quantiles. Similarly, this work makes use of the same quantiles as they provide a good coverage. Additionally, the 5% and 95% quantiles are included as well for added insights in the lowest and highest quantiles.

3.5.3 Data Restrictions

Two more restrictions are made to the data set used to experiment with. The first restriction is about the minimal number of patents a company owns. During this research, only companies that have at least ten patents feature in the data set. This decision is made because of the independent variables that make use of the patents within the patent portfolios of companies: if a company holds too few patents in its portfolio, it becomes difficult to extract information from it. Therefore, only data of companies is kept if they have at least ten patents within their patent portfolio.

The second restriction has to do with the share of patents a company holds within the six technology fields. Since this research is focusing on focal patents from six technology fields, only companies that have at least 50 percent of all patents they own within the six technology fields are present in the data set. This is done to avoid a mismatch between the dependent and independent variables, since the independent variables are calculated by using patents from six technology fields whereas ROA can be affected by patents from other fields as well.

The total number of patents each company holds (also including patents from outside of the six technology fields) is found in the NBER data set mentioned earlier. However, this total number of patents could only be found for the last year that was present in the NBER data set, which was the year 2006. Therefore, for every company, the total number of patents it holds in the six technology fields in the last available year was compared to the total number of patents found in the NBER data set. For example, if company A has 13 patents from the six technology fields in the year 1996 (the last year of which there is data available for company A), but has 30 patents in total according to the NBER data set, the company data would not have been used in this research since $\frac{13}{30} * 100\% \approx 43\% < 50\%$.

Robustness tests have been performed with both the minimal number of patents and the minimal patent share, these robustness tests are discussed in the Results section.

4 Results

This chapter contains the results of the experiments. For all experiments, the focus lies on observing the effects of disruptiveness mean, variance and skewness of the patent portfolio distribution. First, the data set (including the patent citation network it is based upon) used in the experiments is described in detail before discussing the results from the fixed effects OLS regressions, after which fixed effects quantile regression results are shown before ending with a comment on the robustness of the experiments.

4.1 Patent Citation Network

The patent citation network on which the novelty measure was applied is a result of combining the six different technology fields mentioned in Section 3.2.1. Statistics of each individual citation network as well as the citation network obtained from combining all of them are shown in Table 3. In terms of the number of nodes and edges, the citation networks of the medical technology and basic materials chemistry fields are the largest, followed by the networks of the organic fine industry and biotechnology fields while the analysis of biological materials field results in the smallest citation network.

The number of dead ends was found by counting the number of nodes that have an outdegree of zero. For each citation network, the number of dead ends is equal to approximately half of all the nodes in the network. However, the number of dead ends should not influence the calculations of the novelty measure for focal patents much since, for every focal patent, all of the patents it cites are included regardless of the technology fields of the cited patents. Besides, for every technology field, all patents that had links from all other technology fields (beyond the six fields included in this research) were included in the networks as well.

In terms of the average indegrees and outdegrees, the differences between the networks appears to be small. Similarly, the average clustering coefficient of each network is almost equal to zero, indicating that nodes within the networks do not tend to cluster together. A reason for this can be that cycles do not occur within the network because new patents can only cite old patents that were filed before them, while it is impossible for old patents to cite patents that were filed after the older patent's filing date.

Lastly, the degree assortativity coefficients for all networks show that there is a weak correlation between nodes with similar degrees (in cases where the coefficient is positive) and different degrees (in cases where the coefficient is negative). So, within these networks, there is no evidence that nodes tend to create links with other nodes because they are similar or different in terms of degree.

Table 3: Citation Network Statistics

	AOBM	Med	OFI	Bio	Pharm	BMC	All
Nodes	267.562	796.584	682.210	323.505	438.208	836.474	1.982.774
Edges	765.875	3.371.961	1.913.599	1.073.478	1.471.689	2.237.307	7.658.576
Dead Ends	119.669	352.867	337.001	136.867	207.404	427.180	960.160
Avg. Indegree	2,86	4,23	2,80	3,32	3,36	2,67	3,86
Avg. Outdegree	2,86	4,23	2,80	3,32	3,36	2,67	3,86
Avg. Cluster. Coeff.	0,02	0,03	0,03	0,03	0,03	0,02	0,03
Degree Assortativity	0,02	0,15	0,03	-0,002	0,06	-0,03	0,10

An overview of statistics of the patent citation networks for technology fields 11 - Analysis Of Biological Materials, 13 - Medical Technology, 14 - Organic Fine Industry, 15 - Biotechnology, 16 - pharmaceuticals and 19 - Basic Materials Chemistry as well as a large patent citations network used in this research that consists of the other six patent citations networks.

4.2 Descriptive Statistics

The analysis below is performed on the entire data set that is used in the experiments. Keep in mind that this data set includes only data of companies that have at least ten patents in their patent portfolios and have a patent share of at least 0.5 (see Section 3.5.3). The analysis consists of an overview of the companies and observations present in the data set, some descriptive statistics on the variables, a correlation matrix and a look at the Variance Inflation Factors (VIFs).

4.2.1 Companies And Observations

To analyze the data set, companies are linked to the technology field in which they are mostly active. Each company is assigned to one of the six technology fields used in this research, which is done by looking at the field in which the company holds the most patents within the data set. In case of a tie, a random choice between the specific technology fields is made.

For example, if company A has ten patents in total, of which seven belong to the pharmaceuticals field and the rest belongs to the biotechnology field, company A will be assigned to the pharmaceuticals field. If company A were to have five patents that belong to the pharmaceuticals field and five patents that belong to the biotechnology field, company A is linked to one of these fields according to a random choice.

The data set includes 415 different companies and holds 3414 observations in total. Most observations and companies belong to the pharmaceuticals field, while the least number of observations and companies belong to the analysis of biological materials field. This can be observed in Table 4.

Table 4: Companies And Observations Per Technology Field

Technology Field	Companies	Observations
11 - Analysis Of Biological Materials	12	94
13 - Medical Technology	137	1178
14 - Organic Fine Industry	24	175
15 - Biotechnology	44	237
16 - Pharmaceuticals	180	1506
19 - Basic Materials Chemistry	18	224
Combined	415	3414

The total number of companies and observations in the data set per technology field.

4.2.2 Variables

Some additional descriptive statistics on the variables are given in Table 5. The table includes information about the mean, standard deviation, minimum and maximum values as well as the 25, 50 and 75 percentiles for every variable.

Table 5: Descriptive Statistics

	Mean	Std.	Min	25%	50%	75%	Max
Patents (ln)	4.51	0.98	3.71	3.87	4.13	4.65	8.49
Average Citations (ln)	3.61	0.15	3.43	3.50	3.57	3.67	4.55
Total Citations (ln)	5.56	1.40	3.43	4.51	5.27	6.28	10.8
ROA	-0.10	0.53	-17.77	-0.25	0.04	0.17	0.63
Disr. Mean†	0.07	0.07	-0.35	0.02	0.06	0.10	0.45
Disr. Variance†	0.03	0.03	0.00	0.01	0.02	0.04	0.32
Disr. Skewness†	1.82	1.53	-7.02	1.11	1.85	2.73	7.18
Employees (ln)	3.56	0.29	3.43	3.44	3.44	3.51	5.03
Patents† (ln)	4.49	0.95	3.53	3.79	4.03	4.53	8.48
Avg. Citations† (ln)	3.63	0.15	3.43	3.52	3.60	3.70	4.55

Descriptive statistics for all variables in the data set. This table shows the mean, standard deviation (std.), minimal and maximal values as well as the values at the 25, 50 and 75 percentiles in order to show the dispersion and shape of the distribution of each variable. A † indicates that a variable represents statistics based on patents from the previous three years.

4.2.3 Correlation Matrix

Figure 3 shows the correlation matrix for all variables in the data set. Notice here that there are two pairs of variables that show a strong positive correlation and that there is one pair of variables that shows a moderately strong correlation. First of all, total citations and the number of patents show a strong positive correlation. This seems to be logical as having a lot of patents in the portfolio is likely to increase the overall number of citations that portfolio will receive.

Second of all, the dependent variable average citations and the independent lagged variable average citations† show a very strong correlation because they are very similar, with the difference being that average citations† holds information on the number of average citations received over the past three years, while average citations holds data on the number of average citations in the current year.

Third of all, there is a moderately strong positive correlation between patents† and employees. A possible explanation for this is can be that if companies have more employees, it would allow them to let more people work on inventions resulting in more patents.

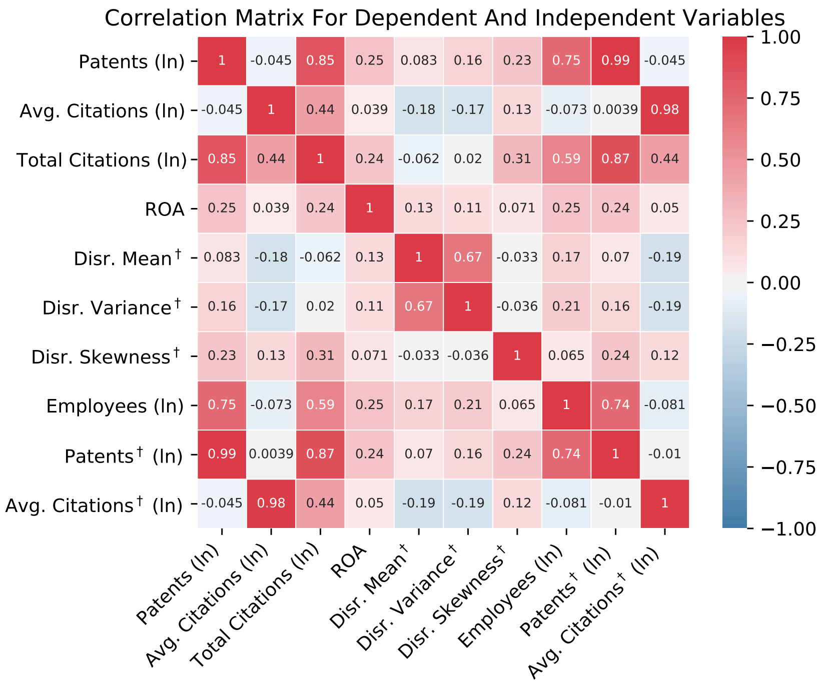

Figure 3: A heatmap showing the correlations between the variables used in this research. The colors vary from warm (red, positive correlation) to cold (blue, negative correlation). A † indicates that a variable represents statistics based on patents from the previous three years.

4.2.4 Variance Inflation Factors

The Variance Inflation Factors (VIFs) for the independent variables are presented in Table 6. As a rule of thumb, a VIF at least 10 is considered to be high, while a value of at leas 5 can also be used. Nevertheless, the VIFs are relatively low and not higher than 2.47, indicating that multicollinearity between the predictors should not be a problem when conducting the experiments.

Table 6: Variance Inflation Factors

Variable	VIF
Disr. Mean†	1.88
Disr. Variance†	1.89
Disr. Skewness†	1.11
Employees (ln)	2.39
Patents† (ln)	2.47
Avg. Citations† (ln)	1.06

An overview of the calculated Variance Inflation Factors (VIFs) for the independent variables. A † indicates that a variable represents statistics based on patents from the previous three years.

4.3 OLS Regression

The results of the OLS regressions are stated below where the order of results is determined by the dependent variables. First, the models that incorporated company fixed effects are discussed, after which a light is shed on the models that did not incorporate company fixed effects. Within each of these sections, results for the patents dependent variable are shown first, the results for average citations and total citations are shown second and third, respectively, and results for ROA are discussed last.

4.3.1 Models With Company Fixed Effects

A summary of the results of the models is shown in Table 7. This table contains the results of four models (one model for each dependent variable) by showing the regression coefficients with standard errors next to them in parentheses. The models incorporate time and company fixed effects and are discussed below.

Patents

Studying the number of patents a company owns can give insights in its productivity. Table 7 shows a negative relationship between the disruptiveness mean and the number of patents: as the disruptiveness mean increases by 1, the number of patents will decrease by 0.70%. This means that having more disruptive patents on average results owning less patents in total and, therefore, in less productivity.

This can support the reasoning given for the differences in the relationships between disruptiveness mean and the two dependent variables of total citations and average citations.

Furthermore, there exists a significant positive relationship between disruptiveness variance and the total number of patents: more disruptiveness variance is linked to a higher productivity. An increase of disruptiveness variance by 1 results in an increase in the number of patents of 0.78%. This can be related to the positive relationship between disruptiveness variance and innovative performance. As stated earlier, it is better for firm performance to have a varied patent portfolio in terms of disruptiveness. Moreover, the diversification of a patent portfolio is achieved by coming up with more innovations that vary in how disruptive they are. So, a company must be productive to diversify its current patent portfolio further, which may explain the positive relationship between disruptiveness variance and productivity.

Average Citations
Average citations can give a good impression of the quality of patents within a patent portfolio. The results in Table 7 show that the coefficients for the variance variable are insignificant, while the other coefficients are significant.

Interestingly, the relationship between disruptiveness mean and total citations is negative, while the relationship between disruptiveness mean and average citations is positive. This means that an increase of disruptiveness mean by 1 results in a 0.07% increase in average citations. So, having more disruptive patents on average leads to more average citations, but it also results in less total citations. An explanation for this observation may lie in the fact that it is more difficult to come up with an invention that is disruptive than one that is stabilizing. Therefore, disruptive innovation leads to less patents and less total citations. However, disruptive innovations have a higher value/impact/quality than stabilizing patents, which explains the positive relationship between disruptiveness mean and average citations.

Furthermore, similar to the model with total citations as dependent variable, there exists a positive relationship between disruptiveness skewness and average citations, meaning that a patent portfolio that is skewed right and holds more stabilizing than disruptive patents is likely to result in more citations on average.

Total Citations
The number of citations patents within a company is used as an additional measure of a firm's innovative performance. Table 7 shows that all coefficients were significant for this model.

The disruptiveness mean variable has a significant negative coefficient, meaning that there is a significant negative relationship between the disruptiveness mean and the number of citations. This means that having a higher average disruptiveness within the patent portfolio leads to worse company performance. Specifically, as the average disruptiveness increases by 1, the total number of citations will decrease by 1.24%. This might imply that, when a company conducts a lot of disruptive innovations while conducting relatively little stabilizing innovations, it may not be able to capitalize on the disruptive innovations and develop useful solutions out of them. This is in line with previous studies that highlight the risks of incorporating too much disruptive innovation. However, bare in mind that the disruptiveness measure is only modestly correlated with impact when trying to interpret the patent portfolio distributions [Funk and Owen-Smith, 2016].

Furthermore, both the disruptiveness variance and skewness variables show a significant positive relationship with total citations. In particular, an increase of 1 in disruptiveness variance results in an increase in total citations of 2.25%, meaning that a higher disruptiveness variance leads to a better innovative firm performance. Literature already suggested that having a mix of both exploring and exploiting patents will result in better firm performance [Geerts et al., 2017] [Gupta et al., 2006]. Moreover, this is related to imbalances in the patent portfolio in terms of disruptive and stabilizing patents leading to worse overall performance as mentioned earlier.

The positive relationship between disruptiveness skewness and total citations is relatively weak: an increase in disruptiveness skewness of 1 leads to an increase in innovative performance of 0.04%. This is means that skewed right patent distributions (where there are more stabilizing than disruptive patents) contribute to more total citations within companies. So, results imply that having a small amount of strongly disruptive patents while having a large amount of strongly stabilizing patents at the same time leads to improved innovative performance. This observation is important as it tells something about the composition of patent portfolios, whereas previous works only mentioned that having both destabilizing and stabilizing patents is important or talked about the average degree of exploration within patent portfolios.

ROA

The summary shows that four of the coefficients are insignificant. This, combined with the model's R^2 score of 0.01 indicates that the predictors were not powerful enough to make good predictions for this specific regression model. The average cited variable was significant with $p < 0.05$ and its coefficient shows a positive relationship between the number of average citations within the patent portfolios over time and ROA. This means that, as patents in the patent portfolio within a company get cited more, company performance increases as well. This reasoning seems logical, as it can be likely for a company that performs well to own patents that get cited more on average.

Table 7: A Summary Of OLS Regression Results With Company F.E.

	Patents (ln)	Avg. Citations (ln)	Total Citations (ln)	ROA
Disr. Mean†	-0.70 (0.10)*	0.07 (0.02)*	-1.24 (0.16)*	0.05 (0.38)
Disr. Variance†	0.78 (0.20)*	-0.07 (0.05)	2.25 (0.35)*	-0.46 (0.38)
Disr. Skewness†	0.00 (0.00)	0.01 (0.00)*	0.04 (0.01)*	0.00 (0.01)
Employees (ln)	0.23 (0.04)*	-0.05 (0.01)*	-0.28 (0.00)*	0.07 (0.04)*
Patents (ln)†	0.81 (0.01)*	-0.01 (0.00)*	1.13 (0.02)*	0.04 (0.03)
Avg. Citations (ln)†	0.00 (0.03)	0.88 (0.03)*	2.73 (0.11)*	0.50 (0.08)*
Invercept	0.18 (0.19)	0.62 (0.12)*	-8.36 (0.52)*	-2.35 (0.37)*
Year F.E.	Yes	Yes	Yes	Yes
Company F.E.	Yes	Yes	Yes	Yes
Companies	415	415	415	415
Obs	3414	3414	3414	3414
Log-likelihood	4734.10	8913.00	2781.40	-931.12
$\mathbf{R^2}$	0.87	0.87	0.80	0.01

A summary of the results for the three panel OLS regression models, one model with both company and time fixed effects for every dependent variable (Patents, average citations, total citations and ROA). For every model, the coefficients, standard errors and p-values are shown (the standard errors are reported in parentheses next to the coefficients). Below the dashed line, for every model, the number of companies and company-year pairs (Obs) in the data set are shown, as well as which fixed effects are incorporated and the log-likelihood and R^2 values. A * indicates that the p-value is significant ($p < 0.05$), while ** indicates that the p-value is less significant ($p < 0.10$). A † indicates that a variable represents statistics based on patents from the previous three years.

4.3.2 Models Without Company Fixed Effects

The same experiments as the ones mentioned above are also conducted without incorporating company fixed effects into the models. A summary of these results is shown in Table 8. This table contains the results of four models (one model for each dependent variable) by showing the regression coefficients and standard errors next to them in parentheses. The models all incorporate time fixed effects but no company fixed effects, their outcome is discussed below.

Patents

There do not seem to be differences between the signs of the variables of disruptiveness mean, variance and skewness for the models with and without company fixed effects. For the negative relationship between disruptiveness mean and the number of patents, the same reasoning that was given for the model that incorporates company fixed effects applies.

Average Citations

The models that do not incorporate company fixed effects show less coefficients that are significant than the models that incorporate company fixed effects. Furthermore, similar results between the models can be observed and similar conclusions can be drawn.

Total Citations

When comparing the company fixed effects model with the model without company fixed effects, it becomes clear that the biggest difference in results lies in the disruptiveness variance variable. The coefficient becomes significantly negative when company fixed effects are disabled, while it is significantly positive with company fixed effects enabled. This means that, when looking at between company differences, more variance in the patent portfolio leads to less total citations, while with within company differences, more variance leads to more total citations.

An explanation for this observation may be found in the way companies are managed. When comparing between companies, a higher variance in disruptiveness might reflect poor management and quality control, which can have a negative effect on company performance. However, when looking at management within a company, the quality of management and quality control are relatively fixed and, generally, do not vary much over the years. In this work, disruptiveness variance is more about diversifying the portfolio, which has a positive effect on firm innovative performance.

ROA

Just as with the model that incorporates company fixed effects for ROA, the predictors for this model are not powerful enough. There are, however, more significant coefficients with this model. The most important one is the result for the disruptiveness skewness variable. There is a significant positive relationship between the disruptiveness skewness of a patent portfolio distribution and ROA: an increase in disruptiveness skewness of 1 leads to an increase in ROA of 0.01. This means that companies of which the distribution of the novelty scores within the patent portfolio is skewed right tend to perform better. At the same time, this implies that a company owns more consolidating patents than destabilizing patents.

This effect is observed when looking at between company differences, while it was insignificant when looking at within company differences. It can possibly explained by a lack of within-firm variance for to exploit in the fixed effect models, meaning that it might be difficult for a company to change its patent portfolio in terms of skewness, which may cause the insignificance of the effect when looking at within company differences.

Table 8: A Summary Of OLS Regression Results Without Company F.E.

	Patents (ln)	**Avg. Citations (ln)**	**Total Citations (ln)**	**ROA**
Disr. Mean†	-0.25 (0.04)*	0.06 (0.01)*	-0.41 (0.08)*	0.18 (0.17)
Disr. Variance†	0.09 (0.07)	0.01 (0.01)	-0.58 (0.14)*	0.41 (0.29)
Disr. Skewness†	0.00 (0.00)*	0.00 (0.00)	0.04 (0.00)*	0.01 (0.00)*
Employees (ln)	0.00 (0.01)	0.00 (0.00)	-0.20 (0.03)*	0.20 (0.02)*
Patents† (ln)	1.03 (0.00)*	0.00 (0.00)	1.33 (0.01)*	0.08 (0.01)*
Avg. Citations† (ln)	-0.12 (0.01)*	0.94 (0.01)*	3.89 (0.06)*	0.46 (0.04)*
Invercept	0.43 (0.05)*	0.19 (0.03)*	-13.72 (0.24)*	-2.93 (0.18)*
Year F.E.	Yes	Yes	Yes	Yes
Company F.E.	No	No	No	No
Companies	415	415	415	415
Obs	3414	3414	3414	3414
Log-likelihood	3134.30	7775.10	-3.20	-2471.90
R^2	0.99	0.97	0.97	0.09

A summary of the results for the three panel OLS regression models, one model with time fixed effects but without company fixed effects for every dependent variable (Patents, average citations, total citations and ROA). For every model, the coefficients, standard errors and p-values are shown (the standard errors are reported in parentheses next to the coefficients). Below the dashed line, for every model, the number of companies and company-year pairs (Obs) in the data set are shown, as well as which fixed effects are incorporated and the log-likelihood and R^2 values. A * indicates that the p-value is significant ($p < 0.05$), while ** indicates that the p-value is less significant ($p < 0.10$). A † indicates that a variable represents statistics based on patents from the previous three years.

4.4 Quantile Regression

Quantile regression is performed on the following quantiles: 0.05, 0.1, 0.25, 0.5, 0.75, 0.9 and 0.95. These quantile regression models make use of the same regression setup (including independent variables, company and year fixed effects etc.) as the setup used in the fixed effect OLS regression models shown in Table 7. The results of the quantile regressions are shown in the same order of dependent variables as used previously. For each quantile regression, a coefficient analysis is performed first before analyzing the results predicted by the quantile regression models. In order to analyse the predictions, specific figures are used that may require some explanation. These figures consist of two plots: the first plot (on the left) shows the fitted values of the dependent variable in which each curve represents results for a different quantile for one dependent variable. The second plot (on the right) is similar but shows the fitted value of the dependent variable but with curves that are fixed[1].

4.4.1 Patents

Figure 4 shows a representation of the coefficients from the patents quantile regression model. A summary of the output of the model is given in Table 9. The quantile regression model with patents as dependent variable is not able to offer additional information because there does not seem to be variation across different quantiles. More result figures are given in Appendix B.

4.4.2 Average Citations

Regarding the quantile regression results for the average citations dependent variable, the same reasoning applies. The coefficients for independent variables per quantile from the model with average citations as dependent variable are shown in Figure 5, the corresponding quantile regression summary is shown in Table 10.

All coefficients for disruptiveness mean, variance and skewness proved to be insignificant and, again, no variation is observed across the different quantiles. Additional figures on these results are provided in Appendix C.

4.4.3 Total Citations

Figure 6 shows the coefficients for each independent variable per quantile for the quantile regression model with total citations as dependent variable. A summary of the output of the quantile regression model is shown in Table 11.

The quantile regression results shown in Figure 6 do not offer added insights a lack of variation in results across the different quantiles (see the figures in Appendix D) The OLS regression results discussed previously seem to apply across the entire distribution.

[1]The curves are fixed around their center and have the form $y_q = a_{vq} * x_v + c$ where a equals the regression coefficient of independent variable v for quantile q with fitted value x_v. The curve is fixed by finding an appropriate value for c by solving the equation for (x^*, y^*) where x^* is equal to the mean of all values of v and y^* is equal to quantile value of the actual y-values in the data set.

Table 9: F.E. Quantile Regression Summary With Patents (ln) As Dependent Variable

	5%	10%	25%	50%	75%	90%	95%
Disr. Mean†	0.00 (0.02)	0.00 (0.03)	0.04 (0.05)	0.12 (0.09)	0.06 (0.12)	0.04 (0.23)	-0.19 (0.35)
Disr. Variance†	-0.02 (0.04)	-0.01 (0.05)	-0.04 (0.10)	-0.15 (0.17)	0.06 (0.12)	0.07 (0.54)	0.34 (0.96)
Disr. Skewness†	0.00 (0.00)	0.00 (0.00)**	0.00 (0.00)*	0.00 (0.00)	0.00 (0.00)	0.00 (0.01)	-0.01 (0.01)
Employees (ln)	0.00 (0.00)	0.00 (0.00)	0.02 (0.02)	0.09 (0.04)*	0.08 (0.04)*	0.06 (0.05)	0.05 (0.07)
Patents† (ln)	0.00 (0.00)	1.00 (0.00)*	1.00 (0.00)*	1.00 (0.01)*	1.02 (0.01)*	1.04 (0.02)*	1.04 (0.02)*
Avg. Citations† (ln)	0.02 (0.01)*	-0.04 (0.01)*	-0.07 (0.01)*	-0.15 (0.02)*	-0.24 (0.03)*	-0.32 (0.03)*	-0.41 (0.05)*
Intercept	0.08 (0.03)*	0.16 (0.04)*	0.23 (0.08)*	0.31 (0.13)*	0.64 (0.15)*	1.02 (0.21)*	1.49 (0.32)*

The coefficients (with the p-values in parentheses) of the independent variables for the fixed effects quantile regression model with patents as dependent variable. Note that the regression setup (including independent variables, company and year fixed effects etc.) is the same as the setup used in the fixed effect OLS regression models shown in Table 7. A * indicates that the p-value is significant ($p < 0.05$), while ** indicates that the p-value is less significant ($p < 0.10$). A † indicates that a variable represents statistics based on patents from the previous three years.

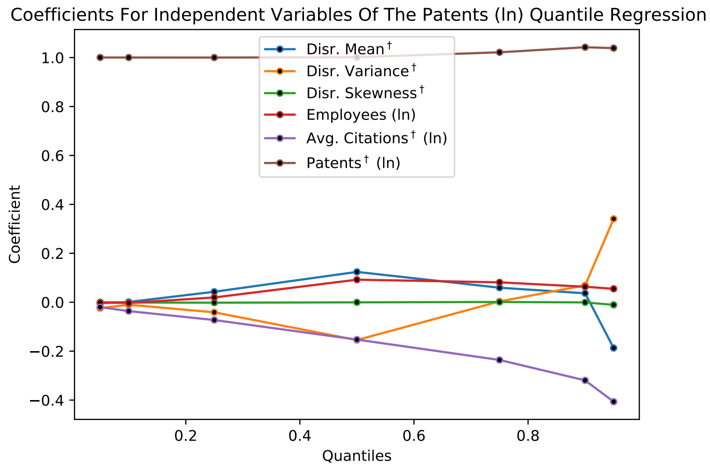

Figure 4: A line graph showing the coefficients of the independent variables that are shown in Table 9 for each quantile. A † indicates that a variable represents statistics based on patents from the previous three years.

Table 10: F.E. Quantile Regression Summary With Avg. Citations (ln) As Dependent Variable

	5%	10%	25%	50%	75%	90%	95%
Disr. Mean†	0.03 (0.03)	0.01 (0.02)	0.00 (0.01)	-0.01 (0.02)	0.00 (0.01)	0.00 (0.00)	0.00 (0.00)
Disr. Variance†	-0.04 (0.07)	0.00 (0.05)	0.01 (0.03)	0.03 (0.03)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)
Disr. Skewness†	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Employees (ln)	-0.01 (0.00)	-0.01 (0.00)	0.00 (0.00)	0.00 (0.00)	0.01 (0.00)	0.00 (0.00)	0.00 (0.00)
Patents† (ln)	0.01 (0.00)*	0.01 (0.00)*	0.00 (0.00)*	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Avg. Citations† (ln)	0.70 (0.03)*	0.77 (0.02)*	0.87 (0.01)*	0.97 (0.01)*	1.01 (0.00)*	1.00 (0.00)*	1.00 (0.00)*
Intercept	1.01 (0.09)*	0.76 (0.08)*	0.41 (0.05)*	0.10 (0.05)*	-0.02 (0.01)*	-0.02 (0.00)*	0.00 (0.00)*

The coefficients (with the p-values in parentheses) of the independent variables for the fixed effects quantile regression model with average citations as dependent variable. Note that the regression setup (including independent variables, company and year fixed effects etc.) is the same as the setup used in the fixed effect OLS regression models shown in Table 7. A * indicates that the p-value is significant ($p < 0.05$), while ** indicates that the p-value is less significant ($p < 0.10$). A † indicates that a variable represents statistics based on patents from the previous three years.

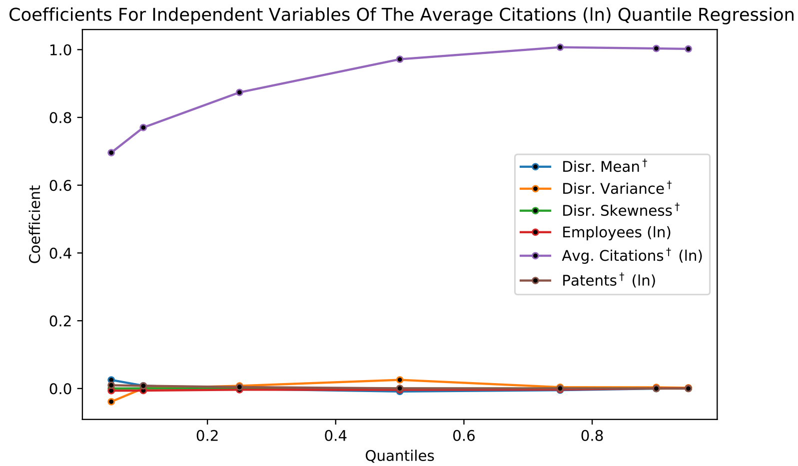

Figure 5: A line graph showing the coefficients of the independent variables that are shown in Table 10 for each quantile. A † indicates that a variable represents statistics based on patents from the previous three years.

Table 11: F.E. Quantile Regression Summary With Total Citations (ln) As Dependent Variable

	5%	**10%**	**25%**	**50%**	**75%**	**90%**	**95%**
Disr. Mean†	-0.58 (0.46)	-0.39 (0.31)	-0.26 (0.17)	-0.29 (0.23)	-0.30 (0.30)	-0.37 (0.26)	-0.45 (0.22)*
Disr. Variance†	-0.22 (0.87)	-0.08 (0.59)	-0.19 (0.32)	-0.35 (0.44)	-0.52 (0.63)	0.02 (0.59)	0.28 (0.46)
Disr. Skewness†	0.02 (0.01)**	0.03 (0.01)*	0.05 (0.01)*	0.05 (0.01)*	0.03 (0.01)*	0.01 (0.01)	0.01 (0.00)
Employees (ln)	-0.06 (0.15)	-0.07 (0.14)	-0.24 (0.11)*	-0.36 (0.11)*	-0.43 (0.13)*	-0.22 (0.18)	-0.02 (0.11)
Patents† (ln)	1.24 (0.05)*	1.26 (0.04)*	1.32 (0.03)*	1.37 (0.04)*	1.43 (0.04)*	1.46 (0.05)*	1.47 (0.05)*
Avg. Citations† (ln)	3.67 (0.26)*	3.91 (0.19)*	4.15 (0.15)*	4.30 (0.17)*	4.39 (0.18)*	4.65 (0.21)*	4.72 (0.22)*
Intercept	-13.41 (1.03)*	-14.27 (0.83)*	-14.62 (0.62)*	-14.82 (0.70)*	-14.89 (0.79)*	-16.65 (1.04)*	-17.51 (0.86)*

The coefficients (with the p-values in parentheses) of the independent variables for the fixed effects quantile regression model with total citations as dependent variable. Note that the regression setup (including independent variables, company and year fixed effects etc.) is the same as the setup used in the fixed effect OLS regression models shown in Table 7. A * indicates that the p-value is significant ($p < 0.05$), while ** indicates that the p-value is less significant ($p < 0.10$). A † indicates that a variable represents statistics based on patents from the previous three years.

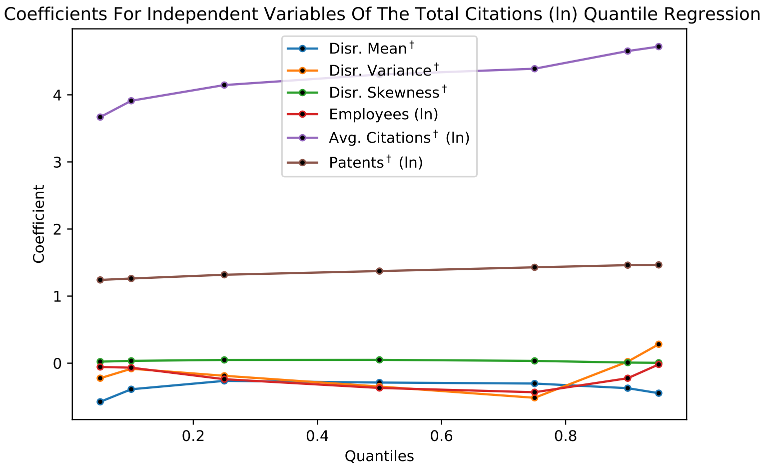

Figure 6: A line graph showing the coefficients of the independent variables that are shown in Table 11 for each quantile. A † indicates that a variable represents statistics based on patents from the previous three years.

4.4.4 ROA

Table 12 shows a summary of the output of the fixed effects quantile regression model for ROA, while Figure 7 shows the value of the coefficients of the independent variables for every quantile used in the quantile regression.

When examining Table 12, it becomes clear that the disruptiveness mean has a significant positive effect on financial performance at the lower quantiles, but not at the higher quantiles. This means that, if a company were to invest more in disruptive innovation on average, it makes sure that the company can achieve a relatively good baseline in terms of minimum performance. However, it does not significantly raise the chance that the company can achieve success to extreme levels. Since a certain level of exploration is needed for a company to survive in competitive markets, this result makes sense.

A better illustration is given in Figure 8. Moreover, it shows that as the disruptiveness mean increases, the range of the dependent variable distribution shrinks, which is driven by raising the lower quantiles. Consequently, a company with a low disruptiveness mean faces a higher level of performance uncertainty, which makes it especially likely to perform among the worst. Furthermore, when the mean disruptiveness is high, the performance uncertainty is low and the company is able to at least stay competitive in the market.

The results show that there are no significant coefficients for disruptiveness variance. The corresponding quantile regression line plots are shown in Appendix E.

Upon further inspection of the disruptiveness skewness variable in Table 12, it becomes evident that it has a significantly positive effect on financial performance at the high quantiles (50% and higher). This may imply that if a company were to own a patent portfolio that is skewed right in terms of disruptiveness, it is likely to result in good maximal financial performance. So, it is better for company performance at the higher quantiles if the company owns more stabilizing patents and a few (strongly) destabilizing patents, which is in line with literature suggesting that having too many destabilizing patents can be detrimental to financial performance.

Table 12: Fixed Effects Quantile Regression Summary With ROA As Dependent Variable

	5%	**10%**	**25%**	**50%**	**75%**	**90%**	**95%**
Disr. Mean†	2.39 (1.35)**	1.95 (0.98)*	1.06 (0.52)*	0.54 (0.25)*	0.16 (0.20)	-0.01 (0.18)	-0.16 (0.27)
Disr. Variance†	-1.04 (2.89)	-0.79 (2.16)	0.23 (1.03)	0.27 (0.51)	0.06 (0.37)	0.11 (0.42)	0.28 (0.60)
Disr. Skewness†	-0.01 (0.03)	0.00 (0.02)	0.01 (0.01)	0.01 (0.01)**	0.01 (0.01)**	0.01 (0.01)*	0.02 (0.01)**
Employees (ln)	0.07 (0.12)	0.09 (0.09)	0.11 (0.06)**	0.19 (0.06)*	0.16 (0.05)*	0.08 (0.06)	0.13 (0.07)**
Patents† (ln)	0.25 (0.04)*	0.19 (0.04)*	0.10 (0.02)*	0.03 (0.01)*	0.00 (0.01)	0.01 (0.01)	0.00 (0.02)
Avg. Citations† (ln)	0.53 (0.21)*	0.27 (0.18)	0.14 (0.11)	0.17 (0.09)**	0.11 (0.07)	0.08 (0.05)	0.02 (0.10)
Intercept	-4.08 (0.79)*	-2.69 (0.67)*	-1.63 (0.47)*	-1.46 (0.43)*	-0.82 (0.31)*	-0.36 (0.26)	-0.25 (0.40)

The coefficients (with the standard errors in parentheses) of the independent variables for the fixed effects quantile regression model with ROA as dependent variable. Note that the regression setup (including independent variables, company and year fixed effects etc.) is the same as the setup used in the fixed effect OLS regression models shown in Table 7. A * indicates that the p-value is significant ($p < 0.05$), while ** indicates that the p-value is less significant ($p < 0.10$). A † indicates that a variable represents statistics based on patents from the previous three years.

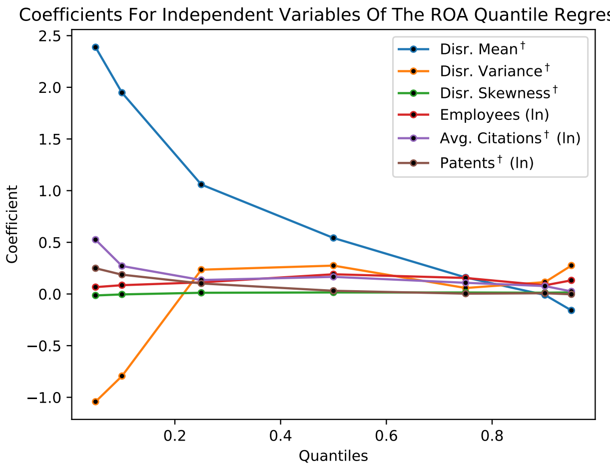

Figure 7: A line graph showing the coefficients of the independent variables that are shown in Table 12 for each quantile. A † indicates that a variable represents statistics based on patents from the previous three years.

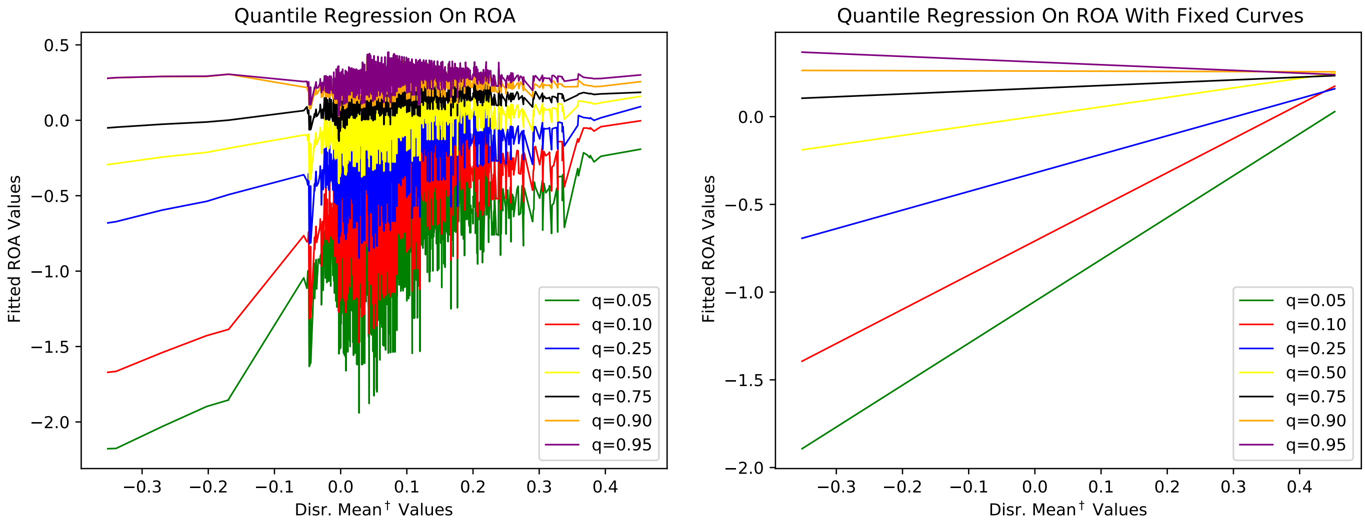

Figure 8: Two line plots that show the predicted ROA values for each quantile based on the mean variable. The first graph (left) shows the predictions that came from the quantile regression model (individual fixed effects incorporated). The second graph (right) shows the predicted ROA values as well, but the curves in this graph are fixed. A † indicates that a variable represents statistics based on patents from the previous three years.

Figure 9 gives more details about this result. Observe here that, for $q \geq 0.50$, the slopes of the lines in the second graph (on the right) are similar and positive. This means that, at the higher quantiles, copmanies are likely to perform better as their disruptiveness skewness increases. This is caused by an increase the amount of stabilizing patents relative to the amount of disruptive patents. Furthermore, due to insignificant results, no definitive conclusions can be made regarding the effect of disruptiveness skewness at the lower quantiles.

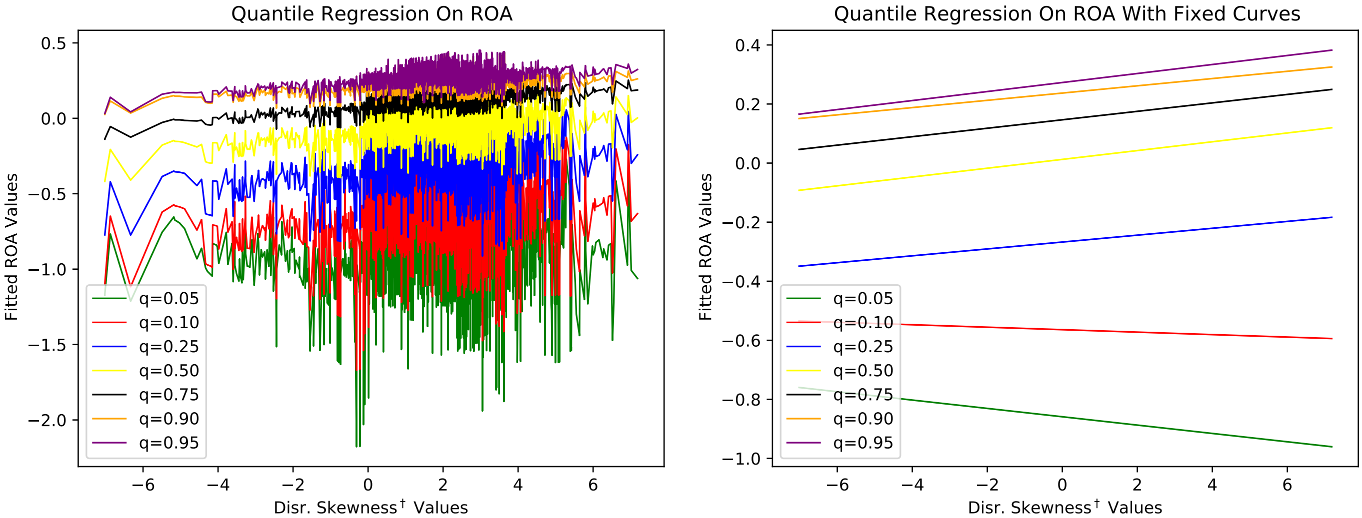

Figure 9: Two line plots that show the predicted ROA values for each quantile based on the skewness independent variable. The first graph (left) shows the predictions that came from the quantile regression model (individual fixed effects incorporated). The second graph (right) shows the predicted ROA values as well, but the curves in this graph are fixed. A † indicates that a variable represents statistics based on patents from the previous three years.

4.5 Robustness Tests

In order to test the robustness of the models shown in this chapter, the experiments were repeated with data sets that varied in the total number of patents a company possesses and the share of patents a company holds within the technology fields used in this work. For the number of patents, experiments were done with a data set including companies with at least 3, 5, 10 and 20 patents. As for the share of patents, experiments were done with a data set including companies with a share of at least 0, 0.25 and 0.5. The results appeared to be robust, except for changing significance of results. However, the coefficients of significant results held the same sign (remained either positive or negative) across experiments.

5 Conclusion

By constructing a data set based on patents from a selection of technology fields, several experiments were conducted that included the use of fixed effect OLS regression models as well as fixed effects quantile regression models. With these experiments, an attempt was made to get a better understanding of how firms could optimize their performance. This method allows for answering the following research question:

What is the optimal distribution of exploring and exploiting patents that will maximize a company's innovative and financial performance?

From the results, it becomes clear that it is better in general to have more exploiting than exploring patents and, in particular, to have a few strongly destabilizing patents while having a relatively large number of stabilizing patents. Furthermore, the company fixed effects models that looked at within company differences showed that a more varied patent portfolio is related to better innovative performance. However, having more disruptive patents leads to less overall innovative productive performance.

Additionally, according to the quantile regression models, companies at the lower quantiles financially benefit the most from having a patent portfolio that is focused on exploring patents. If a company invests more in exploring innovation, it makes sure that the company can hit a good baseline in terms of minimum firm performance. However, this does not significantly raise the chances for companies to achieve extreme success.

Besides, company in the higher quantiles are likely to get a better financial performance with a patent portfolio that is focused on exploitation. However, this is not completely sure due to weaker significance of results, but it is in line with other stating that having more exploiting patents is better for company performance.

So, in some cases, company performance can benefit from a focus on exploring patents, while in other cases, more focus on exploiting patents is preferred. However, attention has to be paid to the success trap and the failure trap, because having an imbalance in the variance of a portfolio can lead to overall worse company performance.

5.1 Theoretical Contributions

This thesis has two theoretical contributions. First of all, the use of quantile regression models has proven to give additional insights to other forms of regression. Similar works have only used regression models to look at average firm performance while not benefiting from the extra layer of data analysis that quantile regression offers. Even though most quantile regression models used in this work showed no difference between quantiles, still one of the models resulted in some significant findings. Therefore, more experimentation should be done with quantile regression models in works that study similar problems.

Second of all, this work has shown that the use of a continuous novelty/disruptiveness measure opens up new research possibilities compared to assuming a dichotomy between exploring and exploiting patents. Previous works have used the dichotomy assumption, but this a relatively simple way of performing patent classification as patents are seldom fully exploiting or exploring. Moreover, the more sophisticated way of classifying patents used in this work allowed for researching the true composition of patents in patent portfolios. This opens up more ways to research different problems that involve analyzing patent portfolios in the exploitation versus exploration debate.

5.2 Policy Implications

This work contributes to the exploitation versus exploration debate and results in two implications for managers and policy makers within companies. Firstly, managers should focus on obtaining a more varied patent portfolio in terms exploring and exploiting patents as this study has shown that is likely to result in improved firm performance. This is in line with previous works stating that both exploration and exploitation should occur simultaneously.

Secondly, managers should generally focus mainly on exploitation, while doing some exploring activities at the same time. This work has shown that owning more exploiting patents while, at the same time, having a few patents that are to a greater extend focused on exploration improves firm performance. However, if a firm performs poorly, managers should focus on exploring activities to obtain a higher level of minimum performance.

5.3 Future Work

There are several ways in which this work can be improved upon. First of all, different technology fields could be used to extend or replace the current data set. It is possible that using different patent citation networks can give different results as these fields have different characteristics. Additionally, some technology fields are more fast moving than others, so incorporating more technology fields might be interesting.

Second of all, the panel OLS regression models with ROA as dependent variable had very low R^2 values. This means that the data and the models do not provide a good fit, which is probably due to the fact that the current predictors were not powerful enough to say something about the ROA dependent variable. So, repeating this research with a different set of predictors might give better insights in financial company performance. Of course, incorporation of different (regression) models might also be considered.

Lastly, it is possible to improve this research by adding weights to patents in the patent citation network used for the disruptiveness calculations. The assumption is made that every patent is as important as any other patent within the patent citation network while, in reality, some patents are of greater importance than others. Therefore, extending the research is possible by taking importance of patents into account during the disruptiveness calculations.

6 References

[Abraham and Moitra, 2001] Abraham, B. P. and Moitra, S. D. (2001). Innovation assessment through patent analysis. *Technovation*, 21(4):245 – 252.

[Alsharkas, 2014] Alsharkas, Z. (2014). Firm size, competition, financing and innovation. *International Journal of Management and Economics*, 44(1):51 – 73.

[Artz et al., 2010] Artz, K. W., Norman, P. M., Hatfield, D. E., and Cardinal, L. B. (2010). A longitudinal study of the impact of r&d, patents, and product innovation on firm performance. *Journal of Product Innovation Management*, 27(5):725–740.

[Baer and Frese, 2003] Baer, M. and Frese, M. (2003). Innovation is not enough: Climates for initiative and psychological safety, process innovations, and firm performance. *Journal of Organizational Behavior*, 24:45 – 68.

[Belderbos et al., 2010] Belderbos, R., Faems, D., Leten, B., and Looy, B. V. (2010). Technological activities and their impact on the financial performance of the firm: Exploitation and exploration within and between firms*. *Journal of Product Innovation Management*, 27(6):869–882.

[Beneito, 2006] Beneito, P. (2006). The innovative performance of in-house and contracted R&D in terms of patents and utility models. *Research Policy*, 35(4):502 – 517.

[Bergek and Bruzelius, 2010] Bergek, A. and Bruzelius, M. (2010). Are patents with multiple inventors from different countries a good indicator of international r&d collaboration? the case of abb. *Research Policy*, 39(10):1321 – 1334.

[Breitzman and Thomas, 2002] Breitzman, A. and Thomas, P. (2002). Using Patent Citation Analysis to Target/Value M&A Candidates. *Research-Technology Management*, 45(5):28–36.

[Breschi and Lissoni, 2005] Breschi, S. and Lissoni, F. (2005). *Knowledge Networks from Patent Data*, pages 613–643. Springer Netherlands, Dordrecht.

[Chang et al., 2009] Chang, S.-B., Lai, K.-K., and Chang, S.-M. (2009). Exploring technology diffusion and classification of business methods: Using the patent citation network. *Technological Forecasting and Social Change*, 76(1):107 – 117. Knowledge Driven Planning Tools for Emerging and Converging Technologies.

[Chen et al., 2011] Chen, X., Wang, X., Wu, D. D., and Zhang, Z. (2011). Analysing firm performance in chinese it industry: Dea malmquist productivity measure. *Int. J. Inf. Technol. Manage.*, 10(1):3–23.

[Choi and Hwang, 2014] Choi, J. and Hwang, Y.-S. (2014). Patent keyword network analysis for improving technology development efficiency. *Technological Forecasting and Social Change*, 83:170 – 182.

[Clarke et al., 2011] Clarke, M., Seng, D., and Whiting, R. (2011). Intellectual capital and firm performance in australia. *Journal of Intellectual Capital*, 12:505–530.

[Coad and Rao, 2008] Coad, A. and Rao, R. (2008). Innovation and firm growth in high-tech sectors: A quantile regression approach. *Research Policy*, 37(4):633 – 648.

[Compustat, 2020] Compustat (2020). Standard and Poor's database for company financials. Data retrieved from Compustat database in May 2020. Data accessed via wrds: https://wrds-www.wharton.upenn.edu/.

[Dabla-Norris et al., 2012] Dabla-Norris, E., Kersting, E. K., and Verdier, G. (2012). Firm productivity, innovation, and financial development. *Southern Economic Journal*, 79(2):422–449.

[Delgado et al., 2002] Delgado, M. A., Fariñas, J. C., and Ruano, S. (2002). Firm productivity and export markets: a non-parametric approach. *Journal of International Economics*, 57(2):397 – 422.

[Ebersberger and Herstad, 2013] Ebersberger, B. and Herstad, S. J. (2013). The relationship between international innovation collaboration, intramural R&D and SMEs' innovation performance: a quantile regression approach. *Applied Economics Letters*, 20(7):626–630.

[EPO, 2020] EPO (2020). European Patent Office. Data retrieved from European Patent Office, PATSTAT Autumn 2019. https://data.epo.org/expert-services/index.html.

[Erdi et al., 2012] Erdi, P., Makovi, K., Somogyvári, Z., Strandburg, K., Tobochnik, J., Volf, P., and Zalányi, L. (2012). Prediction of Emerging Technologies Based on Analysis of the U.S. Patent Citation Network. *Scientometrics*, 95.

[Ernst, 2001] Ernst, H. (2001). Patent applications and subsequent changes of performance: evidence from time-series cross-section analyses on the firm level. *Research Policy*, 30(1):143 – 157.

[Fattori et al., 2003] Fattori, M., Pedrazzi, G., and Turra, R. (2003). Text mining applied to patent mapping: a practical business case. *World Patent Information*, 25(4):335 – 342.

[Fooladi and Chaleshtori, 2011] Fooladi, M. and Chaleshtori, G. (2011). Corporate governance and firm performance.

[Funk and Owen-Smith, 2016] Funk, R. and Owen-Smith, J. (2016). A dynamic network measure of technological change. *Management Science*, 63.

[Furman and Hayes, 2004] Furman, J. L. and Hayes, R. (2004). Catching up or standing still?: National innovative productivity among 'follower' countries, 1978–1999. *Research Policy*, 33(9):1329 – 1354. What do we know Innovation? Selected papers from an International Conference in honour of Keith Pavitt.

[Geerts et al., 2017] Geerts, A., Leten, B., Belderbos, R., and Looy, B. (2017). Does Spatial Ambidexterity Pay Off? On the Benefits of Geographic Proximity Between Technology Exploitation and Exploration: DOES SPATIAL AMBIDEXTERITY PAY OFF? *Journal of Product Innovation Management*.

[Ghani and Ashraf, 2005] Ghani, W. I. and Ashraf, J. (2005). Corporate Governance, Business Group Affiliation, and Firm Performance : Descriptive Evidence from Pakistan. Governance Working Papers 22255, East Asian Bureau of Economic Research.

[Gibson and Birkinshaw, 2004] Gibson, C. and Birkinshaw, J. (2004). The antecedents, consequences, and mediating role of organizational ambidexterity. *Academy of Management Journal*, 47:209–226.

[Grimaldi et al., 2015] Grimaldi, M., Cricelli, L., Giovanni], M. D., and Rogo, F. (2015). The patent portfolio value analysis: A new framework to leverage patent information for strategic technology planning. *Technological Forecasting and Social Change*, 94:286 – 302.

[Gupta et al., 2006] Gupta, A., Smith, K., and Shalley, C. (2006). The interplay between exploration and exploitation. *Academy of Management Journal*, 49.

[Hagedoorn and Cloodt, 2003] Hagedoorn, J. and Cloodt, M. (2003). Measuring innovative performance: is there an advantage in using multiple indicators? *Research Policy*, 32(8):1365 – 1379.

[Han and Sohn, 2014] Han, E. and Sohn, S. (2014). Patent valuation based on text mining and survival analysis. *The Journal of Technology Transfer*, 40.

[Hansson, 2003] Hansson, B. (2003). Company-based determinants of training and the impact of training on company performance: Results from an international hrm survey. *Personnel Review*, 36.

[Huson et al., 2004] Huson, M. R., Malatesta, P. H., and Parrino, R. (2004). Managerial succession and firm performance. *Journal of Financial Economics*, 74(2):237 – 275.

[Jaffe and Trajtenberg, 2002] Jaffe, A. and Trajtenberg, M. (2002). *Patents, Citations, and Innovations: A Window on the Knowledge Economy.*

[Jayaraman et al., 2000] Jayaraman, N., Khorana, A., Nelling, E., and Covin, J. (2000). Ceo founder status and firm financial performance. *Strategic Management Journal*, 21(12):1215–1224.

[Joung and Kim, 2017] Joung, J. and Kim, K. (2017). Monitoring emerging technologies for technology planning using technical keyword based analysis from patent data. *Technological Forecasting and Social Change*, 114:281 – 292.

[Junni et al., 2013] Junni, P., Sarala, R. M., Taras, V., and Tarba, S. Y. (2013). Organizational ambidexterity and performance: A meta-analysis. *Academy of Management Perspectives*, 27(4):299–312.

[Kang and Liu, 2014] Kang, H.-H. and Liu, S.-B. (2014). Corporate social responsibility and corporate performance: a quantile regression approach. *Quality & Quantity*, 48.

[Kim and Bae, 2017] Kim, G. and Bae, J. (2017). A novel approach to forecast promising technology through patent analysis. *Technological Forecasting and Social Change*, 117:228 – 237.

[Koenker, 2004] Koenker, R. (2004). Quantile regression for longitudinal data. *Journal of Multivariate Analysis*, 91(1):74 – 89. Special Issue on Semiparametric and Nonparametric Mixed Models.

[Koski and Mäkinen, 2009] Koski, H. and Mäkinen, I. (2009). Firm size, managerial practices and innovativeness: Some evidence from finnish manufacturing. *International Journal of Technology Management*, 59.

[Kovács, 2019] Kovács, A. (2019). When does it pay off to participate in markets for technology?

[Laforet, 2008] Laforet, S. (2008). Size, strategic, and market orientation affects on innovation. *Journal of Business Research*, 61:753–764.

[Lanjouw and Schankerman, 2001] Lanjouw, J. O. and Schankerman, M. (2001). Characteristics of patent litigation: A window on competition. *The RAND Journal of Economics*, 32(1):129–151.

[Larcker et al., 2013] Larcker, D. F., So, E. C., and Wang, C. C. (2013). Boardroom centrality and firm performance. *Journal of Accounting and Economics*, 55(2):225 – 250.

[Laumann et al., 1983] Laumann, E., Marsden, P., and Prensky, D. (1983). The boundary specification problem in network analysis. *Applied Network Analysis: A Methodological Introduction*, 61.

[Lensberg et al., 2006] Lensberg, T., Eilifsen, A., and McKee, T. E. (2006). Bankruptcy theory development and classification via genetic programming. *European Journal of Operational Research*, 169(2):677 – 697. Feature Cluster on Scatter Search Methods for Optimization.

[Levinthal and March, 1993] Levinthal, D. and March, J. (1993). The myopia of learning. *Strategic Management Journal*, 14:95–112.

[Liu et al., 2011] Liu, S.-H., Liao, H.-L., Pi, S.-M., and Hu, J.-W. (2011). Development of a patent retrieval and analysis platform – a hybrid approach. *Expert Systems with Applications*, 38(6):7864 – 7868.

[March, 1991] March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1):71–87.

[Martinelli, 2010] Martinelli, A. (2010). *The dynamics of technological discontinuities: a patent citation network analysis of telecommunication switches*. PhD thesis, Industrial Engineering & Innovation Sciences.

[NBER, 2020] NBER (2020). National Bureau of Economic Research. Data retrieved from NBER. https://sites.google.com/site/patentdataproject/Home, accessed at April, 2020.

[Noori et al., 2017] Noori, J., Nasrabadi], M. B., Yazdi, N., and Babakhan, A. R. (2017). Innovative performance of iranian knowledge-based firms: Large firms or smes? *Technological Forecasting and Social Change*, 122:179 – 185.

[O'Reilly and Tushman, 2013] O'Reilly, C. and Tushman, M. (2013). Organizational ambidexterity: Past, present and future. *SSRN Electronic Journal*, 27.

[Pakes and Griliches, 1980] Pakes, A. and Griliches, Z. (1980). Patents and r&d at the firm level: A first report. *Economics Letters*, 5(4):377 – 381.

[Powell et al., 1999] Powell, W., Koput, K., Smith-Doerr, L., and Owen-Smith, J. (1999). Network position and firm performance: Organizational returns to collaboration in the biotechnology industry. *Research in the Sociology of Organizations*, 16.

[Purcell, 2019] Purcell, W. (2019). The importance of innovation in business. *Northeastern University Blog*, 24. https://www.northeastern.edu/graduate/blog/importance-of-innovation/, accessed at June 25, 2020.

[Rao et al., 2001] Rao, S., Ahmad, A., Horsman, W., and Kaptein-Russell, P. (2001). The importance of innovation for productivity. *International Productivity Monitor*, 2:11–18.

[Rivettte and Kline, 2000] Rivettte, K. G. and Kline, D. (2000). *Rembrandts in the Attic: Unlocking the Hidden Value of Patents*. Harvard Business School Press, USA.

[Trajtenberg, 1990] Trajtenberg, M. (1990). A penny for your quotes: Patent citations and the value of innovations. *RAND Journal of Economics*, 21:172–187.

[Tushman and Iii, 2006] Tushman, M. and Iii, C. (2006). Ambidextrous organizations: Managing evolutionary and revolutionary change. *Managing Innovation And Change*, 38.

[Uotila et al., 2009] Uotila, J., Maula, M., Keil, T., and Zahra, S. A. (2009). Exploration, Exploitation, and Financial Performance: Analysis of S&P 500 Corporations. *Strategic Management Journal*, 30(2):221–231.

[Verbeek et al., 2004] Verbeek, A., Debackere, K., and Luwel, M. (2004). Science cited in patents: A geographic flow analysis of bibliographic citation patterns in patents. *Scientometrics*, 58(2):241 – 263.

[Verhoeven et al., 2016] Verhoeven, D., Bakker, J., and Veugelers, R. (2016). Measuring technological novelty with patent-based indicators. *Research Policy*, 45(3):707 – 723.

[Wang and Wang, 2012] Wang, Z. and Wang, N. (2012). Knowledge sharing, innovation and firm performance. *Expert Systems with Applications*, 39(10):8899 – 8908.

Appendices

A Relationship Between Field 16 And Other Fields

Table 13 shows the relationships between technology field 16 (pharmaceuticals) and all other technology fields based on the number of patents that were cited by the focal patents that were part of the patent citation network for technology field 16. Note that most citations are going from technology field 16 to patents from the same field. When looking at the top of the table, the technology fields that are related to biotechnology and pharmaceuticals show a strong relation. The weaker the relationships, the more the technology fields get further away in nature from biotechnology and pharmaceuticals, which is to be expected.

Table 13: Overview Of Technology Fields And Their Relationships

Technology Field	Field Name	Number Of Citations
16	Pharmaceuticals	650.173
14	Organic Fine Industry	400.342
15	Biotechnology	278.733
13	Medical Technology	119.625
11	Analysis Of Biological materials	68.882
19	Basic Materials Chemistry	66.609
17	Macromolecular Chemistry, Polymers	63.453
23	Chemical Engineering	46.848
18	Food Chemistry	40.702
29	Other Special Machines	21.732
10	Measurement	15.338
20	Materials, Metallurgy	13.802
28	Textile And Paper Machines	9.750
9	Optics	8.835
21	Surface Technology, Coating	6.349
25	Handling	6.087
34	Other Consumer Goods	4.955
1	Electrical Machinery, Apparatus, Energy	4.498
24	Environmental Technology	4.449
27	Engines, Pumps, Turbines	3.878
26	Machine Tools	3.671
2	Audio-visual Technology	2.729
6	Computer Technology	2.674
31	Mechanical elements	2.090
33	Furniture, Games	2.078
8	Semiconductors	1.978
32	Transport	1.541
30	Thermal Processes And Apparatus	1.279
35	Civil Engineering	1.255
12	Control	1.156
3	Telecommunications	920
5	Basic Communication Processes	421
22	Micro-structural And Nano-technology	408
4	Digital Communication	402
7	IT Methods For Management	253

An overview of all technology fields and their relationship with technology field 16 (including their technology field names) based on the number of times patents from the technology fields got cited by the focal patents from technology field 16. Note that one patent can be assigned to multiple technology fields. This table is an extension of Table 1.

B Quantile Regression Predictions Patents

This appendix contains the graphs belonging to the quantile regression experiments where patents (ln) was used as dependent variable. Figure 10 shows the results for the disruptiveness mean independent variable, Figure 11 show the results for the disruptiveness variance independent variable and Figure 12 show the results for the disruptiveness skewness independent variable. Note that, just as with the models with total citations as dependent variable, the differences between the predicted values for all quantiles are minor.

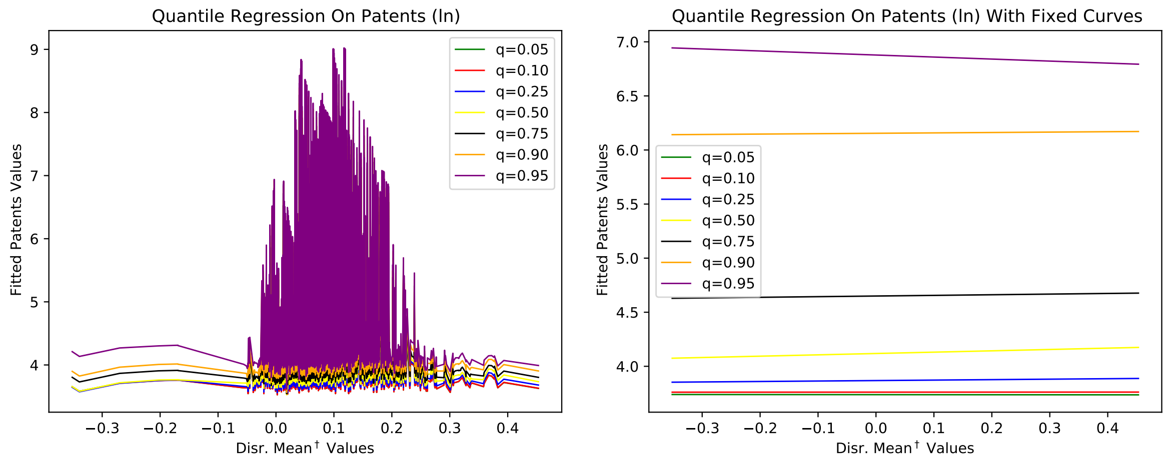

Figure 10: Two line plots that show the predicted Patents (ln) values for each quantile based on the disruptiveness mean† independent variable. The first graph (left) shows the predictions that came from the quantile regression model (individual fixed effects incorporated). The second graph (right) shows the predicted Patents values as well, but the curves in this graph are fixed. A † indicates that a variable represents statistics based on patents from the previous three years.

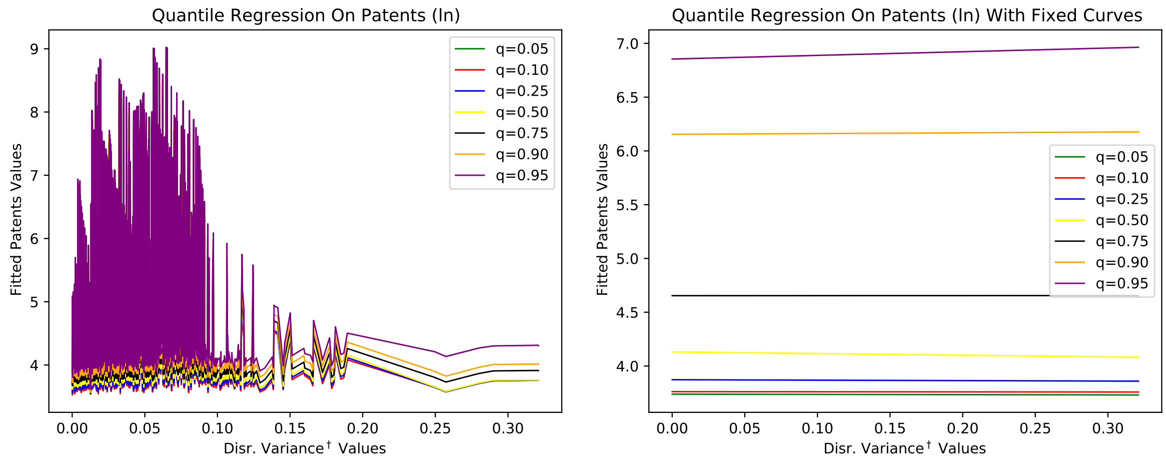

Figure 11: Two line plots that show the predicted Patents (ln) values for each quantile based on the disruptiveness variance† independent variable. The first graph (left) shows the predictions that came from the quantile regression model (individual fixed effects incorporated). The second graph (right) shows the predicted Patents values as well, but the curves in this graph are fixed. A † indicates that a variable represents statistics based on patents from the previous three years.

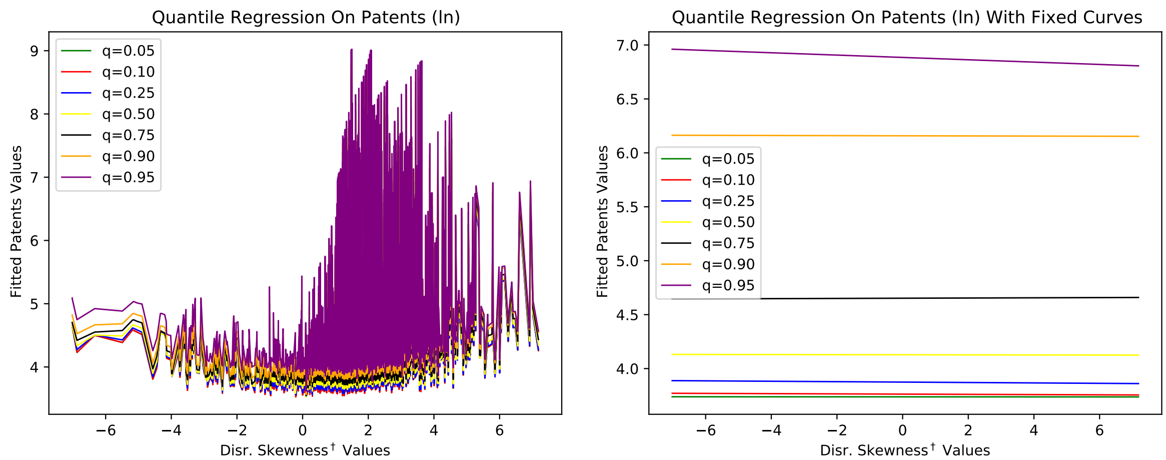

Figure 12: Two line plots that show the predicted Patents (ln) values for each quantile based on the disruptiveness skewness† independent variable. The first graph (left) shows the predictions that came from the quantile regression model (individual fixed effects incorporated). The second graph (right) shows the predicted Patents values as well, but the curves in this graph are fixed. A † indicates that a variable represents statistics based on patents from the previous three years.

C Quantile Regression Predictions Average Citations

This appendix contains the graphs belonging to the quantile regression experiments where average citations (ln) was used as dependent variable. Figure 13 shows the results for the disruptiveness mean independent variable, Figure 14 show the results for the disruptiveness variance independent variable and Figure 15 show the results for the disruptiveness skewness independent variable. Note that, just as with the models with total citations as dependent variable, the differences between the predicted values for all quantiles are minor.

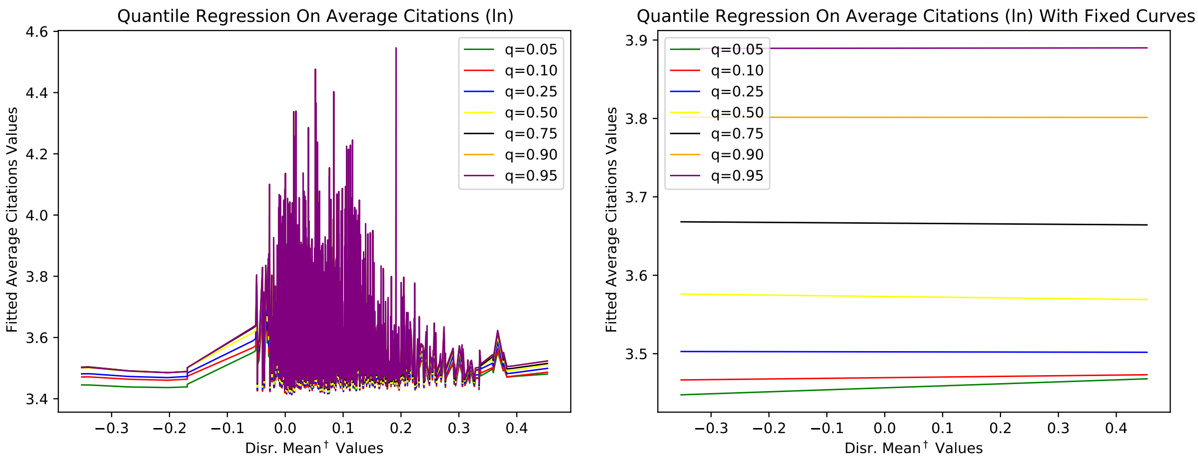

Figure 13: Two line plots that show the predicted Average Citations (ln) values for each quantile based on the disruptiveness mean† independent variable. The first graph (left) shows the predictions that came from the quantile regression model (individual fixed effects incorporated). The second graph (right) shows the predicted Average Citations values as well, but the curves in this graph are fixed. A † indicates that a variable represents statistics based on patents from the previous three years.

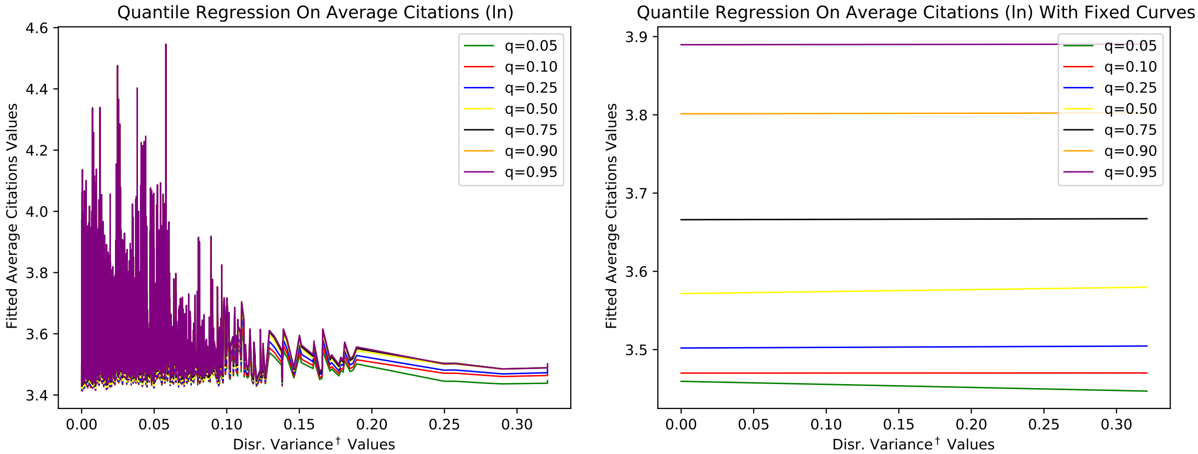

Figure 14: Two line plots that show the predicted Average Citations (ln) values for each quantile based on the disruptiveness variance† independent variable. The first graph (left) shows the predictions that came from the quantile regression model (individual fixed effects incorporated). The second graph (right) shows the predicted Average Citations values as well, but the curves in this graph are fixed. A † indicates that a variable represents statistics based on patents from the previous three years.

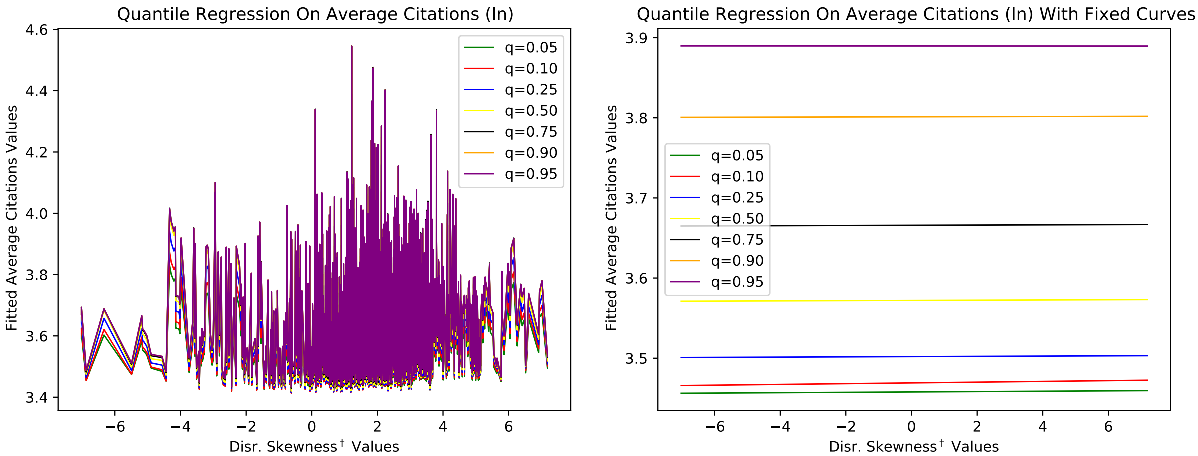

Figure 15: Two line plots that show the predicted Average Citations (ln) values for each quantile based on the disruptiveness skewness† independent variable. The first graph (left) shows the predictions that came from the quantile regression model (individual fixed effects incorporated). The second graph (right) shows the predicted Average Citations values as well, but the curves in this graph are fixed. A † indicates that a variable represents statistics based on patents from the previous three years.

D Quantile Regression Predictions Total Citations

This appendix contains the graphs belonging to the quantile regression experiments where total citations (ln) was used as dependent variable. Figure 16 shows the results for the disruptiveness mean independent variable, Figure 17 show the results for the disruptiveness variance independent variable and Figure 18 show the results for the disruptiveness skewness independent variable. Note that the differences between the predicted values for all quantiles are minor.

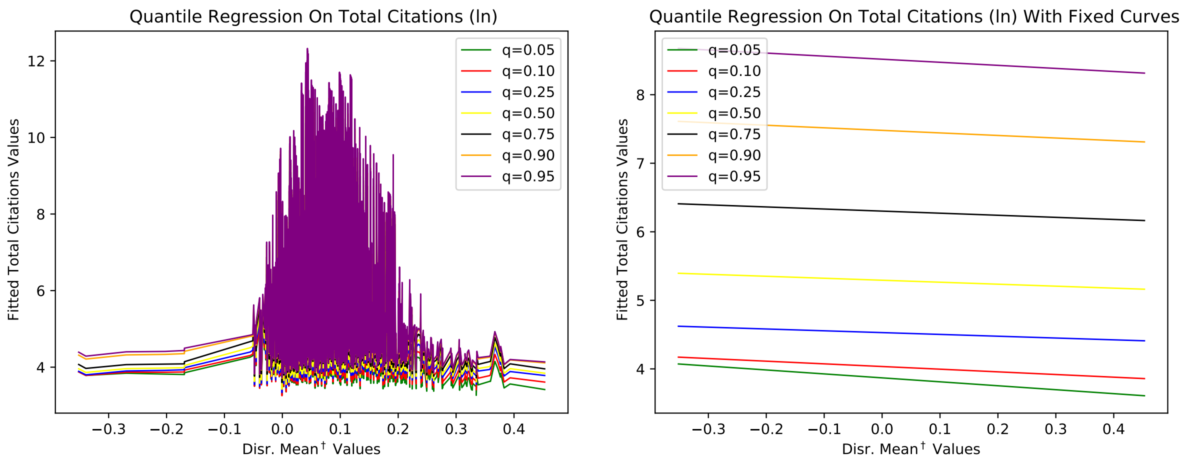

Figure 16: Two line plots that show the predicted Total Citations (ln) values for each quantile based on the disruptiveness mean† independent variable. The first graph (left) shows the predictions that came from the quantile regression model (individual fixed effects incorporated). The second graph (right) shows the predicted Total Citations values as well, but the curves in this graph are fixed. A † indicates that a variable represents statistics based on patents from the previous three years.

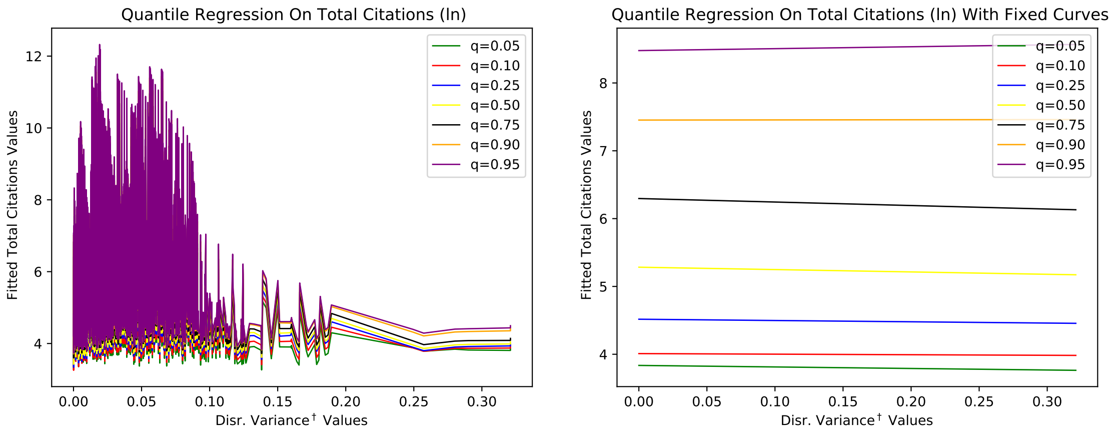

Figure 17: Two line plots that show the predicted Total Citations (ln) values for each quantile based on the disruptiveness variance† independent variable. The first graph (left) shows the predictions that came from the quantile regression model (individual fixed effects incorporated). The second graph (right) shows the predicted Total Citations values as well, but the curves in this graph are fixed. A † indicates that a variable represents statistics based on patents from the previous three years.

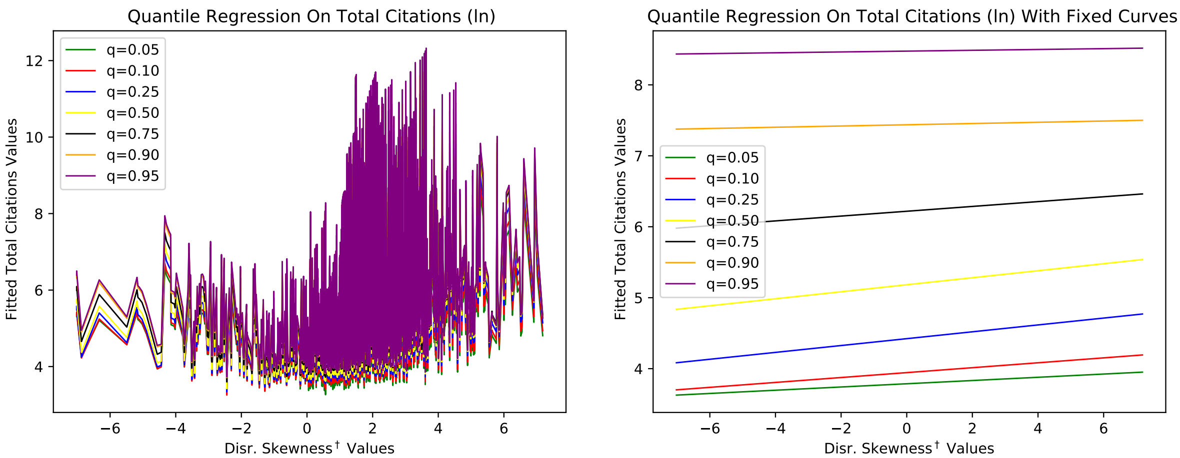

Figure 18: Two line plots that show the predicted Total Citations (ln) values for each quantile based on the disruptiveness skewness† independent variable. The first graph (left) shows the predictions that came from the quantile regression model (individual fixed effects incorporated). The second graph (right) shows the predicted Total Citations values as well, but the curves in this graph are fixed. A † indicates that a variable represents statistics based on patents from the previous three years.

E Quantile Regression Predictions ROA

This appendix contains the graphs belonging to the quantile regression experiments where ROA was used as dependent variable. Figure 19 shows the results for the disruptiveness variance independent variable.

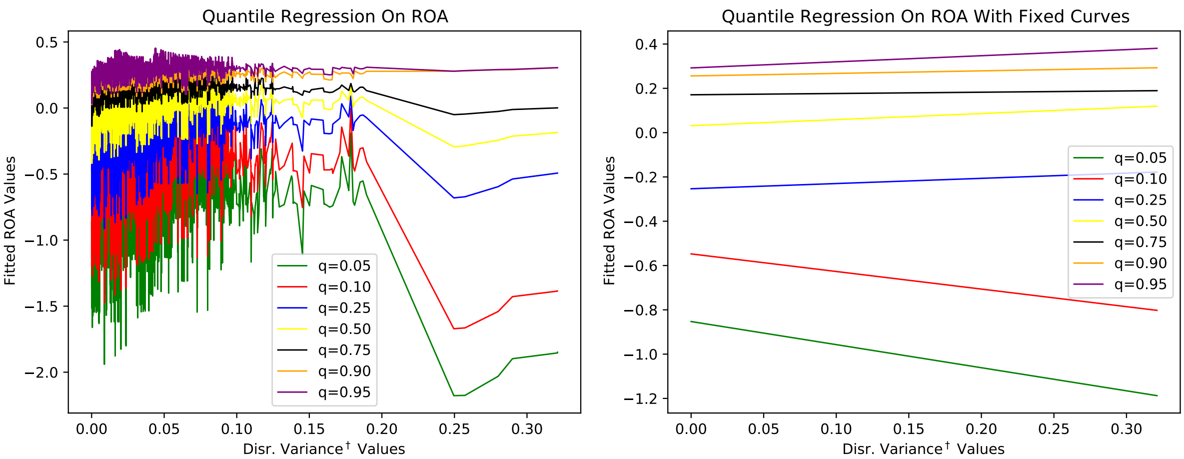

Figure 19: Two line plots that show the predicted ROA values for each quantile based on the disruptiveness variance† variable. The first graph (left) shows the predictions that came from the quantile regression model (individual fixed effects incorporated). The second graph (right) shows the predicted ROA values as well, but the curves in this graph are fixed. A † indicates that a variable represents statistics based on patents from the previous three years.