The Effects of Migration, Competition, and Patents on Innovation

Inaugural-Dissertation
zur Erlangung des Grades
Doctor oeconomiae publicae (Dr. oec. publ.)
an der Ludwig-Maximilians-Universität München

2017

vorgelegt von

Thomas Alberto Fackler

Referent:	Prof. Dr. Florian Englmaier
Korreferent:	Prof. Dr. Klaus M. Schmidt
Promotionsabschlussberatung:	31. Januar 2018

Datum der mündlichen Prüfung: 16. Januar 2018

Namen der Berichterstatter: Florian Englmaier, Klaus M. Schmidt, Oliver Falck

Für meine Familie

Acknowledgments

First and foremost I would like to thank my supervisor Florian Englmaier for his advice, support and encouragement over the last four years. I would also like to thank my second supervisor Klaus Schmidt and Oliver Falck, my third examiner, for their comments, from which my research has greatly benefited.
Part of this thesis has been written at the University of Cambridge thanks to the European Doctoral Group in Economics. I thank Sönje Reiche for this opportunity and I am grateful to Petra Moser for hosting me during a research visit to New York University Stern School of Business. The conversations with them and others at these universities have been very helpful.

I would like to thank my coauthors Yvonne Giesing, Nadzeya Laurentsyeva, Markus Nagler, Monika Schnitzer, and Martin Watzinger. Two chapters of this thesis are based on these joint projects and I have learned a lot through these collaborations.

Conversations with my friends and colleagues in Munich, Cambridge, and New York have benefited my research and made the experience more enjoyable. I thank them and, in particular, my office mates, for discussions after seminars and courses or over a cup of coffee. Moreover, I also thank the colleagues at the chairs of my supervisors for support and for creating such a positive environment.

Financial support from the International Doctoral Program "Evidence-Based Economics" of the Elite Network of Bavaria and from the German Research Foundation (DFG) through CRC TRR 190 is gratefully acknowledged.

Finally, none of this would have been possible without my family. I am grateful for their continuous support and encouragement.

Contents

Preface

The European Union requires member states to promote research and technological development with their industrial policies according to the "Treaty on European Union", which was signed in Maastricht in 1992.[1] China, which has seen unprecedented growth over the last decades as "the world's factory", clearly aims at a smooth transition towards a knowledge-based economy with its current five-year plan for the years from 2016 to 2020: The first of five focus areas is innovation.[2] These are just two examples of technological progress becoming a top priority for policy makers. But why does innovation feature on political agendas with increasing prominence compared to other factors of development and growth?

Economists have long recognized that the accumulation of physical capital cannot explain long-run growth (e.g. Grossman and Helpman, 1994). The technological knowledge in an economy determines its ability to transform labor and capital into valuable outputs. Providing workers with additional factories and machines has diminishing returns, at least if they do not also become better as a result of technological advances. One production factor, whose importance society is less aware of as a decreasing share of the population is employed in agriculture, is fertile soil. As the amount of land the global population can work with is fixed (or may even decline as a consequence of climate change and desertification), the productivity of additional workers is eventually subject to the law of diminishing returns as well. In fact, the first time in human history during which both per capita income and the population have increased over an extended period, has only begun little more than two centuries ago with the Industrial Revolution. Hudson (2014) emphasizes the crucial role of innovation in various areas, extending far beyond steam power and factory production. Already in 1798, the rapid

[1] Treaty on European Union, Title XVII, Article 173

[2] `https://www.uscc.gov/sites/default/files/Research/The%2013th%20Five-Year%20Plan.pdf` (last accessed 7 September 2017).

increase of the population led Malthus to suggest that humanity might be trapped, as food supply could only grow arithmetically, while the population grew geometrically. Any increase in productivity would translate into higher population growth, such that the only stable population size was at subsistence level. It has since turned out that birth rates adjusted eventually and population growth in industrialized countries has slowed down. In some countries, like Germany or Japan, it has even dropped to rates around zero. While parts of the world are still in earlier stages of the demographic transition and the global population continues to grow, current predictions[3] point towards a leveling-off, with a world population of about 11 billion by 2100. At the same time, overall living standards are increasing. Thus, the outlook for humanity is not as gloomy as Malthus' theory predicted. Yet, a stable world population at a higher level requires further technological progress, beyond the developments that enabled such growth in the first place, to limit resource degradation and environmental problems. China's current five-year plan, for instance, recognizes the need for "green growth" in the third of its five major objectives.

Hence, innovation is more important than ever and economists can go beyond incorporating the growth of knowledge into their models, passively modeling how engineers and scientists refine processes and develop products. If the creation and diffusion of knowledge are not exogenous, but emerge endogenously in the economic system, understanding their determinants is crucial (e.g. Romer, 1990; Aghion and Howitt, 1992).

Too many areas of economics have contributed to the identification of the factors that determine technological progress to provide a comprehensive review here. Instead, the following describes how the chapters of this thesis fit into this context and what they add to the literature. The first chapter exploits transitional provisions in the process of European integration as a natural experiment to identify the effects of free movement of labor on the production and diffusion of knowledge. The second chapter studies the effect of competition on the type of corporate research and development (R&D) in the context of the recent surge in Chinese imports to Europe. The third chapter contributes to the literature on patents and innovation, drawing evidence from U.S. antitrust history to estimate how patents affect follow-on innovation, depending on market structure.

[3] `https://esa.un.org/unpd/wpp/Publications/Files/WPP2017_KeyFindings.pdf` (last accessed 7 September 2017).

While concerns about human capital may seem like a recent political issue, the first patent system in the Republic of Venice was established in 1474 precisely to attract skilled artisans from other regions (Moser, 2013). By contrast, concerns about too much immigration through membership in the EU and, in particular, the labor market consequences were hotly debated ahead of the "Brexit" vote in 2016. The impact of international labor mobility is not only controversial in destination countries, however. Emigration is often associated with a loss of human capital, or "brain drain", that hurts the economies of origin countries and prevents them from catching up. In the first chapter[4], we exploit the introduction of free movement of labor in the EU to identify the causal effects of international labor mobility on knowledge production within source countries as well as knowledge flows to these countries. The introduction of free movement of labor after the accession of new EU member states did not come into effect immediately in 2004 and later. The transition periods in the opening of labor markets to migrants from new EU members were decided individually by old EU countries for specific industries. This creates a suitable natural experiment for the identification of the effects of emigration. The transitional provisions created exogenous variation in the ability of new EU members' citizens to migrate to old EU member states.

In response to Paul Krugman's assertion that knowledge flows do not leave a "paper trail", patent citations have been found to be a suitable approximation (Jaffe et al., 1993). Thus, our study not only takes advantage of the fact that, since the Venetian Statute, patent systems have spread throughout Europe and thereby give us a proxy of the amount of innovation. The references to other patents also allow us to trace knowledge flows to emigrants' origins.

We find that rather than suffering from "brain drain", industries with higher migration outflows increase innovation as measured by patents. Moreover, an analysis of cross-border citations suggests that the loss in human capital is compensated by reverse knowledge flows, such that emigration is even helping source countries to converge towards the technological frontier. Consistent with our interpretation, the effects are larger when restricting the sample to high-skilled migrants.

Our findings contribute to the literature on the relationship between labor mobility and innovation in general (e.g. Bosetti et al., 2015; Kaiser et al., 2015), but also add to research

[4]Chapter 1 is based on joint work with Yvonne Giesing and Nadzeya Laurentsyeva.

on international migration and human capital. In the latter literature, several studies take advantage of changes in the H-1B visa program in the U.S. to study the effects of immigration in destination countries (e.g. Kerr and Lincoln, 2010), whereas Kerr (2008) uses the program to study the effects on innovation and productivity in home countries. One important difference between this program and our setting is that H-1B visas are targeted at skilled labor, while we analyze the effects of free movement laws affecting all skill groups.

Furthermore, our study relates to the literature on the localization of knowledge by identifying migration policy as one way to mitigate the effects of country borders, which have been found to be particularly strong obstacles to knowledge flows (Jaffe et al., 1993; Hu and Jaffe, 2003; Singh and Marx, 2013). Our results can help to inform migration policy.

In the second half of the last century, Europe's economies saw increased trade not only within the European Economic Area, but also with other continents. Imports from the world's most populous country have grown so much that a growing literature studies the effects of this "China Shock" on the U.S. and Europe (e.g. Autor et al., 2013; Dauth et al., 2014; Autor et al., 2016b). Economic reforms under Deng Xiaoping, who became the paramount leader of the country in 1978, initiated a transition from socialist command economy to market-based economy. In 2001 China even became a member of World Trade Organization (Naughton, 2007). Building on work by Bloom et al. (2016), I exploit this exogenous shock to study a more general question in the second chapter: How does competition affect firms' R&D strategies? While the relationship between competition and the amount of innovation is an old question that has been investigated at least since Schumpeter (1942), both theoretically and empirically, much less is known about how competition affects the type of R&D that firms invest in. Some theories have been proposed (e.g. Cabral, 2003; Kwon, 2010), but empirical evidence of the causal relationship has been lacking.

Motivated by the theoretical literature on competition and innovation (Aghion et al., 2001), I first explore the effects of competition on firms' investments in an R&D model. The private benefits of large innovations change less with competition than those of small steps, since there is no room for increased business stealing if an innovation is so significant that the innovator becomes a monopolist in any case. This implies a decreasing share of large innovations among all innovations in my model. This relationship holds independently of the initial level of competition, by contrast to the effect of competition on the overall amount of innovation. Yet, this finding is compatible with the inverted-U relationship between competition and total innovation found in the literature (Aghion et al., 2005). Finally, the model shows that the level of competition that maximizes the total amount of large innovations in an industry is lower than the one that maximizes total innovation, thereby providing a theoretical explanation for the trend towards more incremental R&D observed by Arora et al. (2015). Therefore the optimal competition policy changes if the externalities of different types of innovation differ, e.g. because of higher spillovers to other industries for larger innovations.

Technological progress has also affected the methods used in this chapter. Statistical tools for text analysis have been available for a long time (e.g. Deerwester et al., 1990), but only the increase in computational power available to researchers has enabled the recent increase in applications of machine learning methods to large data sets in general and text data in particular (Gentzkow et al., 2017). The empirical part of this chapter develops new measures to quantify the type of innovation. The measures compare patent abstracts of a firm to those of a comparison group such as previous patents of the same firm or the industry. The idea is that a larger step leads to larger changes in the vocabulary in abstracts as well. An advantage of these measures compared to citation-based measures is that they are fixed at the time of application. The number of citations to a patent, by contrast, is determined later on and depends on market conditions itself, while the similarity measures cannot be affected by later developments and, in particular, changes in the private value of inventions as a result of competition.

In instrumental variables regressions, I find support for the theoretical prediction of a decline in the share of large innovations as competition increases due to the growth of Chinese imports. Confirming the model, the empirical results suggest that the relative private value of incremental innovations indeed increases when imports rise.

As this distortion towards smaller innovations should be taken into account in welfare considerations, the empirical estimation of this causal relationship is relevant to various policies, e.g. with respect to trade, standardization, and mergers.

When studying innovation in general, incentives are a natural starting point for economists and intellectual property rights are supposed to provide them for the production of knowledge.

As ideas can be considered public goods, there may be underprovision without intervention. Intellectual property protection thus is an attempt to strike a balance between the inefficiency of a monopoly with incentives for innovation. This tradeoff has been described formally by Nordhaus (1969), but the idea that a temporary monopoly can serve as a reward can already be found in Adam Smith's "Wealth of Nations" (1776).[5] On top of this balance between monopoly and incentives, however, the fact that knowledge is both an input in knowledge production as well as the output has to be taken into account. Any researcher using Google Scholar, the search engine for academic publications, is greeted by the quote "Stand on the shoulders of giants", a metaphor which is often attributed to Isaac Newton, but has been expressed long before. Hence the cumulative nature of science has long been appreciated and economic theory, too, has considered this complication, e.g. to study the effects of patents on follow-on innovators' incentives (Scotchmer, 1991; Green and Scotchmer, 1995). The extent to which such effects are economically relevant, rather than an insight of purely academic interest, is ultimately an empirical question that researchers have begun to address more recently (e.g. Galasso and Schankerman, 2015b; Sampat and Williams, 2015).

The government not only takes an active role in the market by granting temporary monopolies to innovators through the patent system, however. At the same time, the government implements competition policy to prevent firms from abusing their market power. Some remedies available to antitrust authorities have the potential to affect an economy's capacity to produce innovations and need to be studied from this perspective as well. The third chapter[6] of this thesis contributes to the literature by studying one of the most innovative organizations of the 20th century, the Bell Laboratories (Bell Labs), and the consent decree that resulted from an antitrust case against the parent company AT&T in 1956. This ruling has made thousands of patents available without royalties to any company or inventor in the U.S., including those protecting various groundbreaking innovations such as the transistor. Based on forward citations to the affected patents, we estimate that follow-on innovation increased by 17% in the following five years. Young and small companies seem to be the primary beneficiaries

[5] On rewarding merchants for taking risks, Smith writes: "It is the easiest and most natural way in which the state can recompense them for hazarding a dangerous and expensive experiment, of which the public is afterwards to reap the benefit. A temporary monopoly of this kind may be vindicated, upon the same principles upon which a like monopoly of a new machine is granted to its inventor, and that of a new book to its author." (Book V, Chapter I, Part III)

[6] Chapter 3 is based on joint work with Martin Watzinger, Markus Nagler, and Monika Schnitzer.

of the ruling. This suggests that patents pose a particularly strong barrier for these firms, which is consistent with the literature (Galasso and Schankerman, 2015a). Moreover, we find that compulsory licensing can only be an effective antitrust remedy if market foreclosure is addressed as well. In the telecommunications industry, where Bell kept its dominant market position, we do not find an effect on follow-on innovation.[7] Our results on long-run innovation confirm that compulsory licensing increased patenting by about 25% in the affected technology areas.

The chapters of this thesis are self-contained and can be read independently.

[7]Note that this effect of competition policy on the citations to earlier patents confirms the need for the text-based measure in the second chapter, since the novelty of an innovation cannot be affected by later policy changes, but forward citations can.

Chapter 1

Knowledge Remittances: Does Emigration Foster Innovation?*

1.1 Introduction

Remittances, the money international migrant workers are sending back from the country of employment to their home country, represent an important source of income for developing countries and hence constitute a direct benefit of emigration (Russell, 1986). Furthermore, apart from financial contributions, skilled migrants can "send" back the knowledge they have acquired while working in other countries. This remittance of knowledge has the potential to increase innovation in the origin countries and bring them closer to the technology frontier, thus mitigating the negative effects of the loss of human capital due to emigration.

The number of highly educated foreigners in the OECD area now exceeds 31 million, accounting for 45% of the increase in the foreign born population over the last decade (OECD Database on Immigrants in OECD Countries, 2016). The number of skilled migrants has especially increased within Europe since many members of the European Union (EU) and the European Free Trade Association (EFTA) have introduced free movement for citizens of the partner countries. Given the strong increase in labor mobility and rising concerns in countries experiencing net outflows of skilled people, it is important to understand the consequences

*This chapter is based on joint work with Yvonne Giesing and Nadzeya Laurentsyeva.

of migration. Should firms and policy makers think and act in the context of a "global war for talent" or can the international mobility of skilled individuals make everyone better off, in particular, by stimulating cross-border knowledge flows?

In this study, we establish a causal link between labor mobility, knowledge flows, and innovation activities. By exploiting changes in the European labor mobility legislation as a quasi-experimental setting, we evaluate the effect of skilled emigration on innovation. We find that the emigration of skilled individuals increases patenting in source countries and argue that knowledge remittances can explain this positive effect. Using data on patent citations and migration flows from 32 European countries, we find that emigration increases cross-border knowledge flows. Industries that are exposed to a higher mobility of their workers start to cite patents from the emigrants' destinations more frequently than before. The international mobility of skilled workers seems to enlarge R&D networks and promote the transfer of tacit knowledge. In this way, migration enables a faster diffusion of knowledge from more to less technologically advanced countries and helps the latter to catch up.

We embed these results within the following conceptual framework. We assume a knowledge production function, where innovation (here, for instance, measured by the number of patents) is produced with the inputs of capital and labor and a certain production technology. Emigration leads to a reduction in labor and thus has a direct negative effect on innovation production. However, there might also be an indirect effect, which has often been overlooked in this discussion. International migration can increase the flow of ideas and knowledge across borders. Migrants might share knowledge about new technologies, processes, and products with their former colleagues and friends at home. This increases the stock of knowledge in the source countries and, through the recombination of ideas, positively affects innovation. The production technology thus improves and patent production can grow even if the available skilled labor is reduced. Our conceptual considerations thus suggest that migration has a negative direct and a positive indirect effect on patenting levels in source countries. Although we cannot disentangle these effects with our data, we provide empirical evidence on the total effect.

The main challenge in the empirical analysis is the endogeneity of migration flows. This could be due to reverse causality or omitted variables. To establish causality, we construct an instrumental variable (IV) for migration, using changes in labor mobility laws within Europe. These laws are adopted and enforced by the destination countries and hence can be treated as exogenous to economic conditions in migrants' source countries.

The aim of our estimations are twofold. Combining several data sources, we do not only establish a link between emigration and innovation in the source country, but also shed light on the effect of knowledge remittances, potentially driving innovation. We start by analyzing the effects of international labor mobility on total patenting activity in source countries. The IV estimate suggests that a 1% increase in the number of emigrants increases patent applications by 0.64% in the following two years. This result is statistically significant at the 1% level and robust to controls, fixed effects, and varying lags. The effect is quantitatively more pronounced when we consider only the flows of migrants with patenting potential.

We complement the analysis of innovation activity by looking at the convergence in patenting between migrants' origin and destination industries. We limit the sample to pairs where the destination is more technologically advanced than the origin and analyze whether the difference in patenting levels changes with migration flows. This is a highly policy-relevant question, especially in the context of the European Union: Some countries may block the initiatives aimed at enhancing within-EU labor mobility by arguing that the outflow of skilled people will further augment the asymmetries between richer and poorer member states. Contrary to this argument, though, our results show that patenting differences between origins and destinations decrease in the number of emigrants. Hence, emigration can promote convergence to the innovation level of more advanced economies.

To establish the channel for the positive impact of emigration on innovation, we link emigration to reverse knowledge flows, that is the transfer of knowledge from migrants' destinations back to their origins. While skilled emigrants do not patent in their home country anymore, they can stimulate knowledge and technology diffusion, thus improving the production technology in the origin country. Common to the innovation literature, we use cross-border patent citations as a proxy for knowledge flows. The regression analysis relates the number of citations to a particular destination country with the number of migrants that currently work there. We find evidence that knowledge flows from destination to origin indeed increase with migration: the 2SLS regressions yield an elasticity of knowledge flows to emigration equal to 0.59.

Our project relates to two broad strands of the literature. The first one investigates the effects of labor mobility on innovation. Several papers have established a positive effect of migration on patenting in destination countries. Kerr and Lincoln (2010) and Doran et al. (2015) use random visa allocations to find causal effects for the U.S. Bosetti et al. (2015), Parrotta et al. (2014), Ozgen et al. (2014) and Niebuhr (2010) focus on European countries and establish cultural diversity as one of the main channels to generate new ideas and innovation. The effect of migration on source countries received less attention. Kerr (2008) and Choudhury (2015) find that source countries benefit from knowledge flows and return migration and consequently increase patenting and innovation. Kaiser et al. (2015) provide firm-level evidence by looking at worker mobility within Denmark. They find that hiring new knowledge workers increases a firm's patenting activity. Interestingly, the former employers of these workers also increase patenting, which can be explained by reverse knowledge flows. Braunerhjelm et al. (2015) conduct a similar analysis with a matched employer-employee dataset from Sweden and also show that both the receiving and the sending firms benefit from the mobility of knowledge workers. The effects are stronger for interregional mobility. We contribute to this literature by providing causal evidence that emigration leads to an increase in patenting. We thereby confirm what Kerr (2008) and Choudhury (2015) showed for China and India in a very different context and using another methodology. As we have comparable patenting data for source and destination countries, we can extend this result and show that emigration leads to a catch-up process.

The second strand of the literature analyzes the determinants of knowledge flows. Starting with the seminal contribution by Jaffe et al. (1993), these studies have established that knowledge is localized beyond the effects of agglomeration. Later studies focused on international knowledge spillovers (Jaffe and Trajtenberg, 1999; Hu and Jaffe, 2003), showing that knowledge takes time to cross country borders. Thompson and Fox-Kean (2005) challenge the approach by Jaffe et al. (1993) and point out that intra-national localization effects are not robust to a finer technology classification. However, even with their more conservative estimations, the international localization remains significant. Singh and Marx (2013) investigate whether advances in communication technologies and lower costs of traveling reduce the localization of knowledge over time. While they find evidence for a reduction in the significance of state borders in the U.S., their results show that the effect of international borders has even strengthened over time. Few studies so far analyzed the impact of international migration on cross-border knowledge flows.[1] Kerr (2008), for instance, studies the role of skilled immigrants in the U.S. and finds that immigrants form ethnic scientific networks that enhance the technology transfer to source countries.

We extend this literature on knowledge flows to the European context using an identification strategy that allows for a causal interpretation. We build a unique dataset by merging comparable migration data for 32 European countries with European patent data and find evidence for knowledge flows. Due to our unique European enlargement setting, we are able to estimate causal effects of labor mobility independently of other integration events by exploiting different opening times for trade, FDI and migration. We find that the positive effect of mobility on knowledge remittances is particularly high for migrants with patenting potential and is robust to a variety of specifications and samples.

The chapter is organized as follows. The next section describes a conceptual framework to guide our empirical analysis. Section 1.3 outlines the data, followed by Section 1.4 that presents the empirical specification and describes the instrument. Section 1.5 discusses the results. Section 1.6 suggests knowledge flows as the channel. Section 1.7 provides robustness checks and Section 1.8 concludes.

1.2 Conceptual Considerations

This chapter analyzes the effects of emigration on innovation in source countries. As there are two opposing effects, our storyline becomes clearer if we support it with some conceptual considerations. The considerations are based on a classical knowledge production function as introduced by Griliches (1979) and further developed by Jaffe (1986) and Jaffe (1989). We augment the knowledge production function with emigration. The concept illustrates two

[1] Prior literature on international knowledge flows has focused on trade, foreign direct investment and R&D accessibility (Peri, 2005; MacGarvie, 2005, 2006).

opposing effects: a reduction in knowledge production due to a decreasing skilled labor force vs. an increase due to a better production technology induced by knowledge flows and technological spillovers.

We assume a simplified knowledge production function of the form

$$Y = Af(K, L_s). \tag{1.1}$$

K is a measure of relevant capital available for research and development such as laboratories and equipment. L_s stands for skilled labor and A measures total factor productivity (efficiency of knowledge production). In our case A describes how well labor and capital can be combined to produce the knowledge output Y and captures factors that are not explicitly modeled, such as the knowledge stock on which researchers can build. To measure the output Y, we refer to patents, as is common to the literature.

The direct effect of emigration, in this setting, is a reduction in L_s. Due to the outmigration of skilled people, less workers are available for the production of innovation in the source country. The innovation output Y should thus decrease.

However, there is a second indirect effect of emigration that works through the total factor productivity A. After emigration, workers send back knowledge to their home countries. For instance they may transmit technological information and ideas back to their previous employer through communication with former colleagues. This employer becomes better at producing innovation, which is reflected in an increasing A.

Theoretically, it is unclear whether the negative direct or the positive indirect effect prevails. This depends on several other characteristics such as the industry, the technology, and the innovation process. Consequently, it is even more important to gain this knowledge from a rigorous empirical assessment of the question. Using patent data as a measure of innovation output Y and controlling for various other factors corresponding to K and components of A that are unrelated to the stock of knowledge, our empirical specification is able to identify this net effect.

1.3 Data Description

We create a unique dataset by merging comparable migration data for 32 European countries with European patent data. The dataset has four dimensions: origin region[2], destination country, industry (two-digit, NACE Rev. 2), and year. The dependent variables of interest are the number of patent applications (by origin-industry-year) as a proxy for innovation and the number of cross-border citations (by origin-destination-industry-year) as a proxy for knowledge flows. The main explanatory variable is the annual number of emigrants from a given origin currently employed in a given destination industry.

The ideal migration dataset would contain precise data on migration flows, disaggregated by origin and destination (countries and employing industries), skill level, and occupation. In the absence of such a dataset, we use the second-best data from Eurostat Labour Force Surveys (2000 - 2014). These are harmonized surveys, which take place annually in all EU countries, Iceland, Norway and Switzerland and cover around 5% of national populations. The surveys provide demographic information on individuals, including their current country of residence, region of origin (EU15+4, NMS10, NMS2 or Other), education level, occupation, and currently employing industry.[3] We thus obtain the stock of migrants by year, region of origin, destination country, and destination industry. In addition, we can use the information by education level (university degree, vocational degree, or below) and by occupation (two-digit, ISCO) to identify the stock of migrants with patenting potential.[4] The available dataset has several limitations. We can only observe the region of migrants' origin instead of the country. This means that we cannot differentiate between different 2004 accession countries but have to treat them as one region (NMS10). Similarly we have to treat Romania and Bulgaria

[2]Here and in the following text "region" refers to the region to which Eurostat's LFS data aggregate migrants' origin countries: EU15+4 (EU15 and EFTA), NMS10 (new member states in 2004), NMS3 (Bulgaria, Romania and Croatia) and all other countries. The fact that the EU3 region consists of Bulgaria and Romania, which joined the EU in 2007, and Croatia, which followed only in 2013, adds further imprecision, as we cannot tell from the data how many emigrants from this region came from Bulgaria and Romania and were able to take advantage of the EU's right to free movement already.)

[3]EU15+4 include 15 pre-2004 EU member countries (Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, Netherlands, Portugal, Spain, Sweden, United Kingdom) + 4 EFTA countries (Iceland, Liechtenstein, Norway, Switzerland). NMS10 include countries that joined the EU in 2004 (Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Slovakia, Slovenia, Malta, and Cyprus) and NMS2 include countries that joined the EU in 2007 (Bulgaria and Romania).

[4]We assign a dummy called *patenting potential* to migrants working in occupations "Managers" and "Professionals" (ISCO codes: 11, 13, 21, 22, 23, 25, 31, 32, 35).

as one region (NMS2). Furthermore, as we do not observe the origin industry of a migrant, we assume that it is the same as the current industry at the destination. Besides, we cannot identify flows of return migrants. These limitations result in high observational noise and might bias our estimations towards zero.

To construct the instrument for migration flows we use changes in the European labor mobility legislation. We obtain the relevant information from the Labour Reforms database, prepared by the European Commission, which we complement with information from national legislations of the destination countries. Our baseline dataset covers the years from 2000 to 2012, this period encompassed several changes to European labor mobility as described in more detail in Subsection 1.4.2.

The data on innovative activity and knowledge flows come from the EPO's Worldwide Patent Statistical Database (PATSTAT, 2014 Autumn Edition).[5] We are able to assign patents to industries (two-digit NACE Rev. 2) via the International Patent Classification (IPC) of patents.[6]. We then aggregate patent applications by country, industry, and year and patent citations by patenting country, cited country, industry, and year. In our dataset, *patenting country* corresponds to the origin country of migrants, while cited country corresponds to their current destination. To assign patents to countries, we use the PATSTAT information about the location of patent inventors and applicants, which are usually the organizations employing the inventors. Since a patent can have several inventors, it may be assigned to multiple countries if it is the result of an international collaboration. In these cases, we assign a share of the patent to each country that is proportional to the share of co-inventors from that country. The causes and consequences of such collaborations have been studied by Kerr and Kerr (2015). Through this assignment of patents to the inventors' countries it is possible to link a patent with the location of all the patents that cite it.

[5]Patents and patent citations are imperfect measures for innovation and knowledge flows and have been criticized for example by Duguet and MacGarvie (2005). Yet, these are the best proxies, which are available over long periods of time and comparable across the countries we study.

[6]In order to assign four-digit IPC classes to industries, we use the concordance table provided by Eurostat in Appendix 1 of the publication "Patent Statistics: Concordance IPC V8 - NACE REV.2", published in October 2014 and last accessed on 21 November 2016.
https://circabc.europa.eu/sd/a/d1475596-1568-408a-9191-426629047e31/2014-10-16-Final%20IPC_NACE2_2014.pdf

Figure 1.1 motivates the subsequent econometric analysis: cross-border patent citations (a proxy for knowledge flows) significantly increase following the introduction of free labor mobility between a pair of countries. This figure mirrors the response of migration flows to changes in labor mobility regulation within Europe (Figure A.1 in the Appendix).

We complement the dataset with several important control variables: bilateral industry-specific FDI flows (provided by Eurostat), GDP and bilateral trade flows (from CEPII). By combining these different data sources, we can draw conclusions about the effects of international migration on patenting in the origin countries and establish reverse knowledge flows as the channel, while controlling for possible fixed and time-varying confounders.

Figure 1.1: Cross-Border Patent Citations, Annual Treatment Effects of Free Labor Mobility

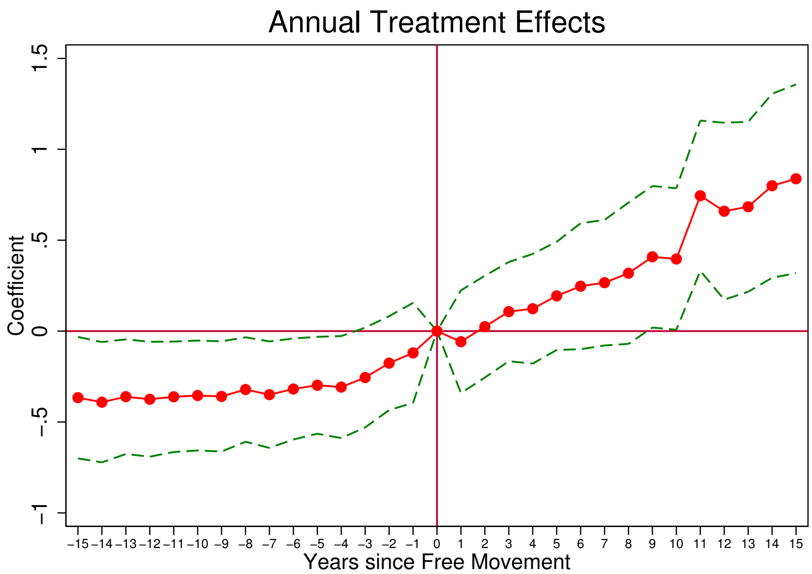

Notes: Annual treatment effects on cross-border citations in patent applications around the introduction of free movement (1965-2014). The regression includes year and country-pair fixed effects. Standard errors are clustered at the citing country level.
Source: PATSTAT, European Commission, own calculations.

1.4 Econometric Specification

In our empirical analysis we first provide causal evidence for the effect of emigration on patenting in origin countries. Second, we link this effect to the increase in knowledge flows. We obtain the elasticities of patenting and cross-border citations to migration using OLS and 2SLS approaches. In the latter, the variation in migration flows is generated only by the exogenous changes in labor mobility laws over 2000-2012. Our baseline regressions include the sample of all patenting European countries. Besides, we provide separate estimates for a sub-sample of Eastern European countries, which were affected the most by the changes in labor mobility over the analyzed period.

1.4.1 Baseline Regressions

Patenting in Origin Countries

We start by analyzing the effects of emigration on total patenting in the origin countries. For this, we aggregate the data at the origin, industry, and year level. Because we do not have detailed country-of-origin data, we use the region of migrants' origin: EU15+4, NMS10 and NMS2. The dependent variable is the number of patent application in a specific origin, industry, and year. The explanatory variable is the number of emigrants from a specific region that work in the same industry but in other European countries. We estimate the following fixed-effects regression:

$$Y_{oiy} = \beta_1 M_{oiy-l} + \beta_2 X_{oiy} + \phi_y + \phi_{oi} + \epsilon_{oiy} \tag{1.2}$$

where o denotes the region of origin, i the two-digit industry, and y the year. Y_{oiy} is the log number of patent applications in a given region and industry. M_{oiy-l} is the log number of emigrants from an origin region, currently working in a given industry.[7] l stands for the lag between migration and patenting. The coefficient β_1 captures the elasticity of patenting to migration. X_{oiy} contains time-varying controls: a dummy for EU membership, trade inflows, and FDI inflows. ϕ_y and ϕ_{oi} denote time and origin-industry fixed effects. ϵ_{oiy} is the error term. The identifying variation thus comes from the within origin-industry changes in the number of emigrants and patent applications. To account for a possible endogeneity bias, we complement the OLS estimations with the 2SLS results, where we instrument migration with changes in labor mobility legislation. We describe the instrument in more detail in Section 1.4.2 below.

[7] Here and in all other specifications, before taking logs we add 1 to each observation. This transformation ensures that we do not lose observations with zero values.

Patenting Asymmetries between More and Less Advanced Countries

We go one step further and analyze whether migration increases or, on the contrary, lowers patenting asymmetries between more and less advanced economies. On the one hand, agglomeration effects and the resources available for research could lead to richer destinations specializing even more on their comparative advantage, thus hindering convergence. If we assume that skilled migrants move from less innovative to more innovative places, labor mobility can increase patenting asymmetries despite some positive effects on the origin. On the other hand, through the migrants working abroad, industries at origins can get access to the frontier knowledge from more advanced economies. This can increase innovation efficiency in origin industries and can allow a faster catch-up process with the technology leaders. Hence, patenting asymmetries between destinations and origins of migrants might decrease. We empirically evaluate the effect of migration on patenting asymmetries with the following regression:

$$log(\frac{P_{diy}}{P_{oiy}}) = \beta_1 M_{odiy-l} + \beta_2 X_{1oy} + \beta_3 X_{2dy} + \beta_4 X_{3odiy} + \phi_y + \phi_{odi} + \epsilon_{odiy} \tag{1.3}$$

The level of observation is origin-destination (od) pair, industry (i), and year (y). The dependent variable $log(\frac{P_{diy}}{P_{oiy}})$ is the log difference in patent applications between the destination and origin industries. The main explanatory variable is M_{odiy-l} - the log number of migrants from origin o working in industry i in destination d. l stands for the lag between migration flows and patenting. The coefficient β_1 shows whether the patenting asymmetries increase or decrease in migration. In this specification we can also control for time-varying origin- and destination-specific effects ($X_{1oy}, X_{2dy}, X_{3odiy}$): the total number of patents at origin, the total number of patents at destination, the total number of patents in a given industry, a within EU dummy (equals one when both origin and destination are EU members), the ratio of GDP per capita between destination and origin, bilateral industry-level FDI, and trade flows. ϕ_y and ϕ_{odi} denote time and origin-destination-industry fixed effects. ϵ_{doiy} is the error term. The coefficient β_1 is thus identified solely through the variation in the number of emigrants within an origin-destination-industry. General changes in patenting at origin and destination cannot confound the results. As with Specification 1.2, we estimate OLS and 2SLS regressions.

Knowledge Flows

Further, we investigate one potential channel behind the effect of migration on innovation: knowledge flows. One speaks of knowledge flows whenever a researcher or an inventor builds on the work done by others to create ideas or to solve a specific technological problem. A common way to track knowledge flows is to use citations data (Jaffe et al., 1993). This approach assumes that a citation to a particular patent or a publication reflects the usefulness of the knowledge contained therein for further work. To determine the effect of migration on knowledge flows we estimate the following empirical model:

$$Y_{odiy} = \beta_1 M_{odiy-l} + \beta_2 X_{1oiy} + \beta_3 X_{2diy-l} + \beta_4 X_{3odiy} + \phi_y + \phi_{odi} + \epsilon_{odiy} \tag{1.4}$$

As in Specification 1.3, the level of observation is origin-destination (od) pair, industry (i), and year (y). The outcome of interest Y_{odiy} represents the log number of cross-border citations. M_{odiy-l} is the log number of migrants from origin o working in industry i at destination d. l stands for the lag between migration flows and patenting. We focus on reverse knowledge flows, i.e., knowledge flowing from destination to origin countries of migrants. Hence, Y_{odiy} represents citations to patents from destination countries by new patents at origin.[8] For example, $Y_{PL/BEiy}$ counts citations by Polish patents in industry i, filed in year y, to existing Belgian patents. It proxies the knowledge flows from Belgium to Poland. $M_{PL/BEiy-l}$ represents the number of Polish migrants in Belgium, currently working in industry i. The coefficient β_1 captures the elasticity of citations to migration. In our example, it shows the percent change in the number of citations from Poland to Belgium if the number of emigrants from Poland to Belgium increased by 1%.

To avoid mechanic effects from the general increase in patenting at origin or destination industries, we control for the number of patent applications in the origin industry (X_{1oiy}) and for the lagged number of patent applications in a destination industry X_{2diy-l}. X_{3odiy} denote other controls: a within EU dummy (equals one when both origin and destination countries are EU members), the total number of patents in a given industry, the bilateral FDI,

[8] We consider citations in patent publications and date patents with their application filing date.

and trade flows. ϕ_y and ϕ_{odi} denote time and origin-destination-industry fixed effects. ϵ_{odiy} is the error term. We again run both OLS and 2SLS regressions.

1.4.2 Instrument for Migration Flows

Even though we control for many observable factors and have a number of fixed effects in the baseline OLS regressions, an endogeneity problem might still arise. Estimates could be biased, for instance, if reduced patenting at the origin forces inventors to leave. To avoid this problem, we use changes in the labor mobility laws in Europe as a source of exogenous variation for migration flows.

The freedom of movement for workers is a policy chapter of the acquis communautaire of the European Union and represents one of the four economic freedoms: free movement of goods, services, labor and capital. According to the Article 45 of the Treaty on the Functioning of the EU, "freedom of movement shall entail the abolition of any discrimination based on nationality between workers of the Member States as regards employment, remuneration and other conditions of work and employment." In practice, it means that there are no restrictions (such as quotas on foreign workers) or additional bureaucratic procedures (such as obtaining a work permit or a permission from the local authorities) related to the employment of foreign citizens. This right primarily concerns the citizens of the EU and EEA member states who, starting from 1958, have gradually introduced free labor mobility towards their partner countries.[9]

In our project, we exploit two episodes of changes in the free labor mobility in Europe. First, in 2004 all EEA countries introduced free movement for the citizens of Switzerland. Switzerland responded with a symmetric measure in 2007.[10] Second, a special scheme has been in force following the EU enlargements in 2004 and 2007. For up to seven years after the accession, old EU members could restrict the access to their labor markets for citizens of new member states. While some countries kept the restrictions for the whole period, some

[9] Norway and Iceland exert this right since 1994. Liechtenstein exerts this right since 1995, but imposes a permanent quota for all EEA citizens.

[10] However, as a result of the "Against mass immigration" initiative, Switzerland is scheduled to impose permanent quotas on residence/work permits for citizens of all EEA countries except Liechtenstein, starting from 2017.

provided easier labor market access only in certain industries, and some opened up their entire labor markets directly upon the accession. When imposing restrictions the countries had to apply them to the whole group of NMS from the same entry year. Therefore, they could not target labor mobility laws at the citizens of some particular states. Iceland, Liechtenstein, Norway, and Switzerland applied the transitional provisions towards the accession countries in the same way. These labor mobility laws created variation in the migration flows between European countries on the country, industry, and year level. Table A.1 in the Appendix provides an overview of the precise opening dates of countries and industries. Importantly for the identification, these changes to labor mobility did not coincide with other integration events (free movement of capital and goods).

Figure 1.2 shows the spikes in migration from NMS during the initial opening in 2004, when countries such as the UK, Sweden, and Ireland opened their labor markets and in 2011 when all transitional provisions for the 2004 accession countries where abolished and Germany, for instance, fully opened its labor market.

We can thus instrument real migration with exogenous labor mobility legislation. The first-stage regression takes the following form:

$$M_{odiy} = \gamma_1 FM_{odiy-1} + \gamma_2 FM_{odiy-2} + \gamma_3 FM_{odiy-3} + \gamma_4 X_{odiy} + \nu_y + \nu_{odi} + u_{odiy} \tag{1.5}$$

FM_{odiy-l} is an indicator variable, which is equal to one if a specific industry i in a destination country d is open for labor migrants from a country o in a given year y. We include a one, two and three year lag to allow for the delayed effect. In our sample this indicator changes only for origin and destination pairs with either Switzerland or new EU member states. As these migration flows might be different, we show separate results for migration from only Eastern Europe in every case. X_{odiy}, ν_y, and ν_{odi} are the same controls and fixed effects as used in the baseline OLS specifications. When using the instrument for the patenting regressions (Specification 1.2), we aggregate the values of the free movement variable by

Figure 1.2: High-skilled Migration in Europe

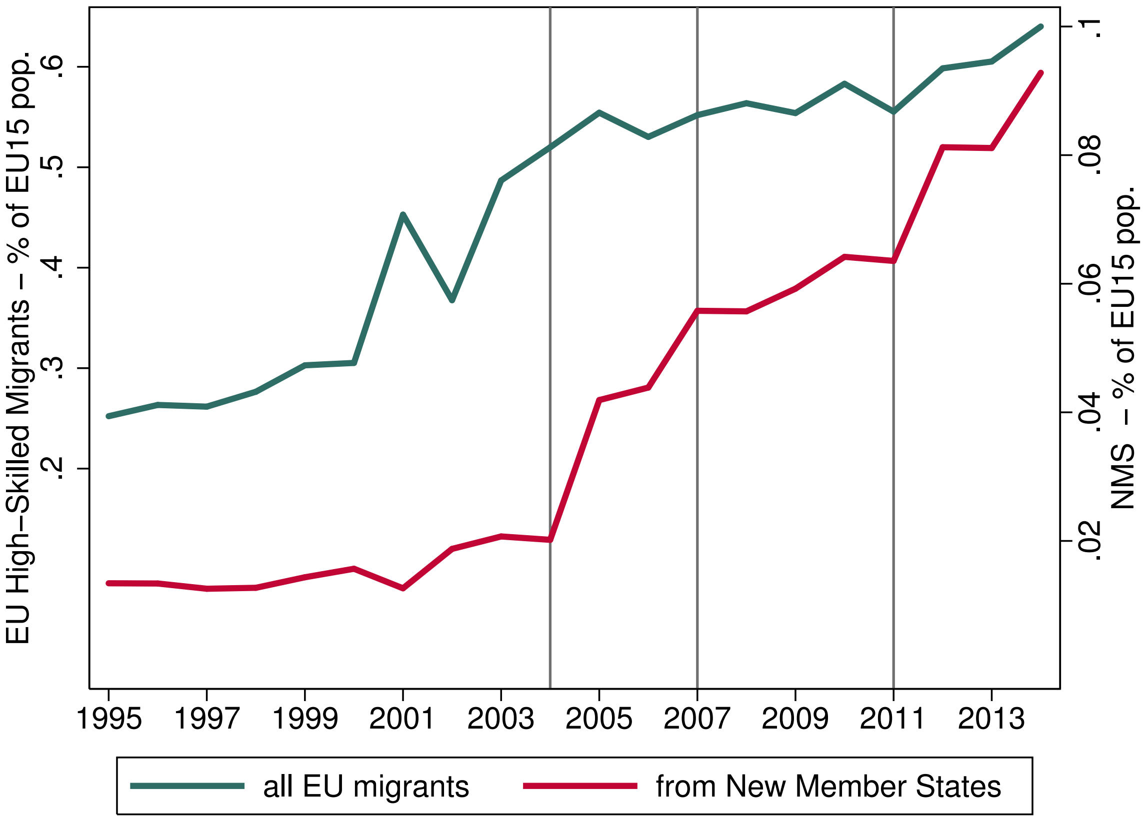

Notes: The graph shows the share of high-skill migrants (born in other European countries) in the EU15 population.
Source: Eurostat.

origin, industry, and year.[11] In this case, the FM variable can be interpreted as the exposure of a given origin-industry (oi) to free labor mobility of its workers.

When constructing the free movement dummies, we take into account the fact that many old EU members did not explicitly specify which industries are open to migrants from the NMS, but rather allowed for special job schemes in sectors that experienced labor shortages. In case of such implicit exceptions, we set the free movement dummy equal to one and multiply it by a measure of labor shortages in a given industry of an old EU member state. As such measure, we use the share of firms (in the destination industries) reporting to be constrained by the factor labor. These data are available from the European Commission Business Survey. To account for possible endogeneity (arising, for instance, when labor shortages are reported in industries that grow faster in all EU countries), we control for the overall number of patent applications in a given two-digit industry (aggregate over all European countries).

1.5 Results

In this section, we first show the effects of migration on total patenting at the origin. Second, we provide evidence that emigration can reduce patenting asymmetries between less and more advanced economies. We show OLS as well as 2SLS results. First-stage and reduced form regressions are provided in the Appendix. Our baseline sample includes all patenting European countries. In addition, we show separate estimations for the sub-sample of Eastern European countries.

1.5.1 Migration and Patenting

This section shows that the emigration of labor increases overall innovation, measured by the number of patent applications per year in a region. As the migration data only allow us to estimate the effect of emigration at the region level, we aggregate the free movement variable by industry and region of origin: EU15+4, NMS10, and NMS2. The aggregated FM

[11] For each origin region we have 31 free movement indicators corresponding to 31 possible destinations. We aggregate them to one measure by using proximity weights (the inverse log distances between the two largest cities of two countries.)

Table 1.1: Patent Applications and Migration, OLS and 2SLS

	(1) OLS Patents	(2) OLS cit. weighted	(3) OLS Patents	(4) 2SLS Patents	(5) 2SLS cit. weighted	(6) 2SLS Patents
L2.Migrants	0.0994***	0.0949**		0.637***	0.903***	
	(0.0259)	(0.0420)		(0.139)	(0.199)	
L2.Migr.pat.potential			0.0572			1.175***
			(0.0420)			(0.332)
in EU	-0.262***	-0.298***	-0.296***	-0.112	-0.0729	-0.406**
	(0.0903)	(0.0752)	(0.0844)	(0.157)	(0.205)	(0.164)
L2.Trade flow	1.634***	2.535***	2.124***	-0.679	-0.945	3.325***
	(0.348)	(0.432)	(0.342)	(0.607)	(0.877)	(0.724)
L2.FDI inflow	2.03e-05**	3.15e-05**	2.10e-05**	1.16e-05	1.84e-05	2.34e-06
	(9.82e-06)	(1.22e-05)	(8.06e-06)	(2.07e-05)	(2.90e-05)	(1.12e-05)
Observations	383	383	383	383	383	383
Region industry FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Clusters	53	53	53	53	53	53
F				6.517	11.29	7.285

Notes: The regressions in this table estimate the relationship between the migration flow out of a country and innovation in that country. The first three columns are estimated with OLS and the last three columns use a 2SLS estimation with our instrument based on free movement legislation. The dependent variables are the number of patent applications in an industry and origin region in a year or, in columns 2 and 5, the citation-weighted patent applications (i.e., patent applications + forward citations to these patents). Patent application numbers and citation-weighted counts, number of migrants and trade flows are taken in natural logarithms (more precisely, for variable x we use $\log(x+1)$ to include observations where $x = 0$). The sample includes all EU members and countries in the European Free Trade Association. All specifications include year and region-industry fixed effects. Robust standard errors are clustered at the region-industry level. Sources: PATSTAT, Eurostat, CEPII *** p<0.01, ** p<0.05, * p<0.1

measure approximates the number of countries to whose labor markets an inventor in a certain industry and region of origin had access and is normalized to be between 0 and 1, where 1 corresponds to full access to all EU15+4 countries.

The first three columns of Table 1.1 show the baseline OLS regressions and the last three columns show 2SLS regressions, which use the labor mobility legislation as an instrument for migration.[12] Columns 1 and 4 estimate the relationship between the overall number of emigrants and the number of patent applications from inventors in that region. These

[12] Note that the right to free movement was not symmetric due to a one-sided transition period, e.g. workers of old EU member states have been able to move to new EU member states as a rule earlier than the other way round. Thus the instrument varies also with the direction of migration and we observe variation in emigration and patenting over time for pairs of origin region and industry. We cluster on the origin-industry level to account for autocorrelation in the regressions in Table 1.1. When we consider asymmetries and citations, there is additional variation depending on the destination country, such that we cluster on the origin-destination-industry level.

regressions show that workers' migration to other EU member states has a significant and positive effect on patenting in the regions of origin. As both variables are measured in logarithms, the coefficient can be interpreted as the elasticity: the effect in the IV estimation in column 4 suggests that a 1% increase in the number of emigrants in an industry causes patent applications in the region of origin to increase by 0.6%. The 95% confidence interval for the elasticity is between 0.37 and 0.91. If we consider the average number of emigrants in the year 2004 (2459 emigrants) and the average number of patent applications 2 years later (255 applications) for new EU member states per industry, this implies that about 1 to 2 additional applications result from 25 additional emigrants.[13] Note however, that this number only includes migrants in industries that were matched to the patent data, i.e., in which there is patenting. Furthermore, the number of patent applications in 2006 we have used for this calculation already includes the additional applications, such that the number of additional patents is likely to be lower. Despite the noise and the level of aggregation in our data the regressions are able to reject that there is a negative effect.

The second and fifth columns of Table 1.1 use citation-weighted patents as the treatment variable, i.e., the number of patent applications plus the number of citations to these patents in a region, industry, and year. The number of later patents building on and therefore citing a patent is often used as a measure of quality.[14] The citations for more recent cohorts in our sample are subject to truncation, which is controlled for through year fixed effects. As the coefficients are similar, we conclude that the quality of patenting has not deteriorated. Thus, merely a higher propensity of inventors in origin regions to file patents as a result of European integration does not seem to be the driver of the effect. Of course, the number of later patents citing a patent (forward citations) is only a rough measure of quality and may be affected by emigrants spreading information about their home countries' latest technologies abroad as well. Nonetheless, a higher number of forward citations would likely be associated with a greater benefit of source countries' innovations, since they indicate that more follow-on innovation built on them.

[13] 1% of 2459 emigrants is about 25 and 0.37% (0.91%) of 255 applications is 0.94 (2.32).

[14] The relationship between citations and the social value of an invention has been documented in a case study on Computed Tomography scanners in Trajtenberg (1990). A more recent study by Moser et al. (2015) finds a robust correlation between citations of hybrid corn patents and the improvement in yield reported in field trial data.

Columns 3 and 6 differ from the other regressions in Table 1.1 in the migration variable, which here includes only emigrants with patenting potential. Whereas the OLS regression shows a smaller and insignificant partial correlation, the coefficient in the IV regression is larger than the corresponding coefficient for all migrants in column 4.

The OLS estimate is likely to be downward biased due to omitted variables and reverse causality. If there is an omitted variable in the OLS regressions that is negatively correlated with emigration and positively with patenting levels, then the OLS estimate is downward biased. This is very likely and could be driven, for instance, by management quality. A good manager might lead to a good work and research environment. This results both in high patenting levels and low emigration from this firm and consequently biases the OLS estimate downward. Moreover, we might encounter reverse causality in the OLS regressions. If higher patenting levels lead to less migration, then we observe a negative relationship between the two variables that goes in the other direction. As a consequence, the OLS estimator is smaller than it should be and thus downward biased.

Tables A.2 and A.3 in the Appendix provide the first stage results and the reduced form that complement the 2SLS results analysis. One can see that the instrument is highly relevant in the first stage and that the overall effect of the three lags for the free movement variables sum up to a positive effect.

Table A.4 in the Appendix provides the same table with the restricted sample of NMS10 countries (2004 accession years). Due to the level of aggregation in the migration data, the 2SLS effects are not significant. Importantly we find no evidence of a significant negative effect, which would be expected if the loss of human capital dominated.

1.5.2 Migration and Convergence

While the results of the previous section suggest that emigration can positively affect innovation at the origin, this section investigates whether this positive effect is enough to reduce patenting asymmetries between less and more advanced economies or whether international migration still benefits knowledge production at destination countries more. This analysis is relevant for policy discussions about benefits and costs of free labor mobility in Europe.

Furthermore, the results in this section serve as a robustness check for the effects found above. When analyzing asymmetries we use all four dimensions of our dataset: origin, destination, industry, and year, and can therefore control for unobserved origin- and destination-specific time-varying changes, which could bias our estimates of patenting elasticity to migration in Section 1.5.1.

To have a clear direction of migration flows from less to more advanced economies, we restrict the sample to the origin-destination pairs, where destinations are EU15+4 countries and origins are new EU member states. In addition, in our baseline sample we consider origin and destination pairs with Switzerland as a destination and other EU15+4 countries as origins. We also show the results for migration from Eastern Europe only, and the results are consistent. For each industry and year, we construct an asymmetry measure as the log difference between the amount of patent applications at destinations and origins.

On average, destination industries file more than three times the amount of patent applications compared to origins. As expected, the patent quality of the former is also higher. We then regress the asymmetry measure on the number of migrants. Table 1.2 presents OLS (columns 1-3) and 2SLS (columns 4-6) results. The OLS coefficient of migration is slightly positive, but is not statistically significant. This may be caused by the bias due to higher migration outflows from more problematic industries. Another reason is that once we move to the more disaggregated level, we introduce more noise in the migration data (more missing and zero observations). This especially concerns already disaggregated migration data by skill and occupation. 2SLS estimates, however, suggest that emigration allows origin industries to catch up to the patenting level of destinations: a 1% increase in the number of migrants reduces patenting asymmetries by 0.30% (column 4 and 5 in 1.2). The coefficient for migrants with patenting potential is much larger in magnitude, but is imprecisely estimated (see column 6). Overall, the regressions' results can be interpreted within the framework of a patent production function with decreasing returns to skilled labor: a marginal increase in patent production at destinations (due to the immigration of skilled labor) is smaller than the marginal increase in patenting at origins (due to the increase in patenting efficiency), e.g. because knowledge about the latest technologies was scarce in source countries, while skilled labor was abundant.

Table A.6 in the Appendix presents the results from the same specifications but estimated on a restricted sample with new EU member states as origins and EU15+4 as destinations (thus excluding emigration from EU15+4 to Switzerland). The obtained coefficients are slightly smaller in magnitude, but still significant. Table A.7 in the Appendix shows the reduced form results, where instead of migration figures we use the bilateral free movement dummies. One of the drawbacks of our migration data is the large amount of missing observations, which could be either due to the effective absence of migrants or to misreporting.[15] This raises external validity issues to our estimations in terms of a generalization to all European countries. Therefore, the most interesting results of Table A.7 are in columns 5 and 6 where we present the coefficients from the regressions over the whole sample of origin and destination pairs. The number of observations increases multiple times, yet the coefficients for the free movement dummies are very close to the estimates from the baseline sample. Moreover, most coefficients are more precisely estimated due to improved power: we note that EU membership, higher bilateral trade flows and FDI also help the convergence.

While interpreting the regression coefficients, we implicitly assume that migrants stay within the same industry. This is reasonable, as for skilled migrants the losses associated with changing the industry are substantial. Hence, they are more likely to seek employment in the same sector in the destination countries. If the assumption would not hold for some industries, how would this affect our estimations?[16] Suppose there are two industries: L and M in Poland and Belgium. The Polish migrants from industry L move to Belgium to work in industry M. Empirically, we observe $M_{BE/PL/M/y}$ to increase. The inflow of the skilled Polish workers in the Belgian industry M raises its innovation activities (or in the worst case, does not affect them). The performance of the Polish industry M is likely to remain unchanged. The asymmetry measure $log(\frac{P_{BE/M/y}}{P_{PL/M/y}})$ either increases or at most stays the same, which goes in the opposite direction of the reported effect. We thus might underestimate the magnitude of the effect.

[15] For example, due to missing migration data we have to drop all observations with Germany as a destination country.

[16] There are pairs of NACE industries, between which inventors may indeed be likely to move, for example between "26 Manufacture of computer, electronic and optical products" and "27 Manufacture of electrical equipment".

Table 1.2: Convergence in Patenting Levels and Migration, OLS and 2SLS

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	OLS	2SLS	2SLS	2SLS
	$log(\frac{P_{diy}}{P_{oiy}})$	$log(\frac{P_{diy}}{P_{oiy}})$	$log(\frac{P_{diy}}{P_{oiy}})$	$log(\frac{P_{diy}}{P_{oiy}})$	$log(\frac{P_{diy}}{P_{oiy}})$	$log(\frac{P_{diy}}{P_{oiy}})$
	Patents	cit. weighted	Patents	Patents	cit. weighted	Patents
L2.Migrants	0.0319	0.0376		-0.305**	-0.334**	
	(0.0223)	(0.0276)		(0.146)	(0.158)	
L2.Migr.pat.potential			0.117**			-1.831
			(0.0575)			(2.212)
Patents, origin	-1.220***	-1.391***	-1.206***	-1.207***	-1.376***	-1.419***
	(0.0762)	(0.0817)	(0.0753)	(0.0883)	(0.0946)	(0.281)
Patents, dest	1.066***	1.105***	1.069***	1.058***	1.096***	1.021***
	(0.0713)	(0.0908)	(0.0717)	(0.0777)	(0.0978)	(0.0894)
Within EU	0.00806	-0.0884*	0.0109	0.0194	-0.0759	-0.0180
	(0.0483)	(0.0531)	(0.0487)	(0.0520)	(0.0572)	(0.0635)
GDP_d/GDP_o	-0.173	0.400	-0.197	-0.188	0.384	0.173
	(0.316)	(0.367)	(0.319)	(0.338)	(0.394)	(0.530)
L3.Trade flow	-0.0791	-0.0236	-0.0718	-0.0281	0.0326	-0.113
	(0.0629)	(0.0799)	(0.0622)	(0.0679)	(0.0867)	(0.0827)
L3.FDI flow	0.000575	-0.000443	0.000380	-0.000116	-0.00120	0.00254
	(0.00668)	(0.00668)	(0.00671)	(0.00783)	(0.00786)	(0.00926)
Observations	2,946	2,946	2,946	2,864	2,864	2,864
R-squared	0.486	0.551	0.486	0.424	0.500	0.325
Origin-dest-industry FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Clusters	582	582	582	500	500	500
F				83.92	122.8	76.50

Notes: The dependent variable is the natural logarithm of $Patents_{dest}/Patents_{origin}$. Number of migrants, number of patents (in origin and destination countries), GDP ratio between destination and origin, FDI, and trade flows are in natural logarithms. The sample includes all EU and EFTA members. All specifications include year and origin-destination-industry fixed effects. Robust standard errors are clustered at the origin-destination-industry level. Sources: PATSTAT, Eurostat, CEPII
*** p<0.01, ** p<0.05, * p<0.1

1.6 The Channel: Knowledge Flows

Having established that emigration leads to an increase in patenting, we want to analyze one potential channel in more detail: knowledge flows. This section shows that migrants stimulate knowledge flows from their new destinations to their countries of origin.

Table 1.3 presents the baseline OLS and 2SLS results. The dependent variable is the log count of citations by patents in the origin to the destination country. This dependent variable proxies the knowledge flows due to emigration. In the baseline estimations, we allow for two-year lags between the time of migration and the citations in the patent applications. The results are similar for a one-year lag but slightly weaker. Importantly, given the structure of the dataset, we can account for origin-industry and destination-industry shocks. A possible threat to identifying the coefficient of interest would arise if destination industries, which experienced a positive patenting shock, started to attract more workers from other countries. A higher supply of patents from this destination would also mechanically increase the amount of citations to this country. We can control for such an effect by including the number of patent applications in the destination industry (with a three year lag).[17] In a similar way, we control for the number of patent applications in the source country. The migration effect is identified from the within origin-destination variation in the migration stocks and the count of cross-border citations. Since both dependent and explanatory variables are in natural logs, the coefficient represents the elasticity of cross-border citations to the number of migrants.

In the first column, we regress the citations on the overall number of migrants M_{odiy}, year, and origin-destination-industry fixed effects; in column 2 we add additional time-varying controls; in column 3 we use the number of migrants with patenting potential as the main independent variable. OLS results suggest a positive association between migration and cross-border citations. The estimated coefficient for migrants with patenting potential is robust to all controls and is twice as large compared to the overall migration stock.

Columns 4 to 6 of Table 1.3 show the 2SLS results that yield quantitatively larger elasticities than the OLS. A 1% increase in emigrants induces a 0.59% growth in cross-border citations to their origins. Table A.8 in the Appendix summarizes the results for the sub-sample where

[17] As a rule of thumb, it takes about three years for a patent to be granted.

Table 1.3: Citations to Destination Industries, OLS and 2SLS

	(1) OLS	(2) OLS	(3) OLS	(4) 2SLS	(5) 2SLS	(6) 2SLS
L2.Migrants	0.0334*	0.0269		0.799***	0.588***	
	(0.0170)	(0.0167)		(0.213)	(0.225)	
L2.Migr.pat.potential			0.0638*			2.916
			(0.0348)			(2.302)
Patents, origin		0.191***	0.192***		0.174***	0.192***
		(0.0237)	(0.0238)		(0.0268)	(0.0310)
L3.Patents, dest		0.0435***	0.0431***		0.0427***	0.0219
		(0.0145)	(0.0145)		(0.0158)	(0.0236)
Within EU		-0.0501	-0.0471		-0.0698*	0.0468
		(0.0378)	(0.0379)		(0.0416)	(0.0845)
L3.Trade flow		0.00665	0.0119		-0.104*	0.00902
		(0.0392)	(0.0390)		(0.0617)	(0.0440)
L3.FDI flow		0.00780	0.00711		0.0126**	-0.0134
		(0.00493)	(0.00495)		(0.00570)	(0.0203)
Observations	7,299	7,287	7,287	7,136	7,124	7,124
R-squared	0.080	0.095	0.095			
Origin-dest-industry FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Clusters	1322	1320	1320	1159	1157	1157
F				20.29	22.20	14.98

Notes: The dependent variable is number of citations from a region and industry to another country per year. Citation counts, number of migrants, total number of patent application in origin and destination industries, FDI and trade flows are taken in natural logarithms. The sample includes all EU and EFTA members. All specifications include year and origin-destination-industry level fixed effects. Robust standard errors are clustered at the origin-destination-industry level. Sources: PATSTAT, Eurostat, CEPII
*** p<0.01, ** p<0.05, * p<0.1

new EU member states are origins and EU15+4 are destinations. Despite the reduction in the sample size, the main 2SLS coefficients remain positive and significant. The reduced form regressions (Table A.9 in the Appendix) are also consistent. When we estimate the reduced form for the whole sample of origins and destinations, the free movement coefficients gain significance and quantitatively remain almost identical to those from the baseline sample. This indicates that some of the insignificant results in the baseline regressions (as, for example, the imprecise coefficient for migrants with patenting potential) are mainly due to power problems with noisy migration data.

Previous research has emphasized the role of communication between moving researchers and their former colleagues at the previous employers (e.g. Kaiser et al., 2015; Braunerhjelm et al., 2015). To test whether the channel they have found for inventors moving between firms within a country is also the primary channel of international knowledge flows in our setting, we exclude the inventor's network. To do this, we exclude citations between inventors and all employers (applicants) and other inventors they are listed with on a patent application at any point in time. Table 1.4 reports the results for the restricted sample. While the coefficients change slightly, they remain positive and significant. Thus only a small part of the effect seems to be driven by the inventors' close network. Knowledge flows that this method could not capture include, for example, if a student at an Eastern European university moves on to work in Western Europe, filing patents for the first time and citing her professors' research. However, the sizable effect that remains suggests that wider spillovers play an important role.

Citations are not always added by the inventor himself but can also be added by the examiner. One worry might thus be that examiners become more aware of research done in other European countries and that they consequently are more likely to add citations from these countries. Alternatively, the effect might be driven by the fact that more patents are filed at the European Patent Office, where examiners may be more likely to add references to foreign patents than at the national offices.[18] This concern is addressed by Table A.11, which shows

[18]The latter concern is also addressed in Table A.10, where only citations among patents filed with the USPTO are included, such that European institutional changes should not affect the results.

Table 1.4: Citations to Inventor's Network Excluded

	(1) OLS	(2) OLS	(3) OLS	(4) 2SLS	(5) 2SLS	(6) 2SLS
L2.Migrants	0.0346**	0.0276*		0.797***	0.631***	
	(0.0171)	(0.0167)		(0.197)	(0.203)	
L2.Migr.pat.potential			0.0464			3.878*
			(0.0314)			(2.355)
Patents, origin		0.174***	0.175***		0.155***	0.177***
		(0.0226)	(0.0226)		(0.0262)	(0.0358)
L3.Patents, dest		0.0353***	0.0350***		0.0344**	0.00732
		(0.0134)	(0.0134)		(0.0150)	(0.0247)
Within EU		-0.0496	-0.0471		-0.0711*	0.0787
		(0.0355)	(0.0356)		(0.0403)	(0.0867)
L3.Trade flow		0.0296	0.0350		-0.0893	0.0312
		(0.0389)	(0.0387)		(0.0597)	(0.0484)
L3.FDI flow		0.00982**	0.00925*		0.0150***	-0.0179
		(0.00481)	(0.00482)		(0.00569)	(0.0224)
Observations	7,299	7,287	7,287	7,136	7,124	7,124
R-squared	0.077	0.091	0.091			
Origin-dest-industry FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Clusters	1322	1320	1320	1159	1157	1157
F				19.18	19.85	10.82

Notes: In this table, citations within the network of the inventor are excluded, i.e., citations from applicants and inventors with whom the cited inventor has patented at any point in time. The dependent variable is the number of citations from a region and industry to another country in a year. Citation counts, number of migrants, total number of patent application in origin and destination industries, FDI and trade flows are taken in natural logarithms. The sample includes all EU and EFTA members. All specifications include year and origin-destination-industry level. Robust standard errors are clustered at the origin-destination-industry level. Sources: PATSTAT, Eurostat, CEPII
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

the results only with citations that were added by the applicant (rather than the examiner or a third party) according to PATSTAT and we can see that there are no qualitative changes.[19]

There are a number of ways for the knowledge flows to occur in practice. One possibility is that emigrants increase the awareness of new knowledge or technologies. This could happen, for example, if emigrants inform their former colleagues or if they meet at conferences. Another possibility is that researchers in the source countries are aware of new knowledge or technologies but need to learn how to use the tacit knowledge embedded in them. A close contact among former colleagues might spur the transfer of tacit knowledge. Additionally

[19] In unreported regressions, we limit citations further to only include those that are marked in PATSTAT as applicant-added and, additionally, where citing and cited patents are both priority patents filed at the USPTO. The results are qualitatively similar despite the fact that only less than 1% of citations remain.

return migration can increase innovation in source countries. Often, emigrants return to their home countries after several years abroad and create start-ups or contribute to innovation in other ways.[20]

1.7 Robustness

To confirm the validity of the results, we conducted a number of robustness checks. We find that the increase in patenting activities as a result of emigration is not driven by different pre-trends or institutional changes in the European patenting system.

One way to check the validity of the results is to examine pre-trends. If our results are valid, the coefficient of interest should be zero if we regress citation patterns on future labor market openings. Figure 1.1 in Section 1.3 shows the annual treatment effects for the regression of cross-border citations on the free movement variable. We look specifically at bilateral citations during the time period 15 years before and 15 years after free movement between two countries has been established. The data we use for this graph are based on patent applications over the 50 year period from 1965 to 2014. The regression includes year dummies and country-pair fixed effects to take out trends. The figure shows that there is no significant change in cross-border citations in the years prior to the establishment of free labor mobility.[21] This is reassuring and increases the credibility of our results. It becomes clear that the effect only starts to gain momentum at the time of the introduction of free movement and builds up over the following years.

One might also worry that the institutional framework of registering patents has changed in the EU, especially in the context of EU enlargement and the European Patent Convention. We thus restrict the sample to patents that have been registered at the United States Patent and Trademark Office (USPTO). Table A.10 in the Appendix shows the results. While we have fewer observations, the qualitative results remain the same. The results thus do not seem to be driven by institutional changes in Europe.

[20] Our time frame of analysis is more likely to reflect the increasing awareness of new technologies or the transfer of tacit knowledge.

[21] Note that this graph uses country-level data, such that the free movement indicator only switches to 1 once all sectors are open. Some of the (insignificant) increase before time 0 may thus be due to the partial openings during the transition periods, which we exploit in the main part of the chapter for identification.

1.8 Conclusion and Policy Implications

This study analyzes the effects of emigration on patenting levels in source countries. We find that countries that experience emigration increase their level of patenting. We further suggest that this has led to a catch-up process that brought origin countries closer to the technology frontier. We also find that the international mobility of people has increased technology and knowledge spillovers as evidenced by cross-border patent citations in the respective countries. Specific channels that could have fostered the knowledge spillovers are the transfer of tacit knowledge, the increased and improved network of inventors and return migration.

One policy recommendation that directly follows from these findings is that the EU could benefit from further facilitating migration within Europe. As there are no more legal barriers to free labor mobility, hindering factors are mostly language and administrative barriers. The EU could reduce these barriers by ensuring the recognition of foreign qualifications and the promotion of language courses at all age levels. In this way, the EU can exploit the full potential of migrants both for destination and source countries.

Another policy implication is to ease skilled migration to Europe from outside the European Union. This could be achieved by easing the access to European labor markets and the recruitment of highly qualified foreign workers. While the Blue Card has been a step in this direction, its scope could be increased to obtain a higher impact and administrative barriers should be reduced. For those skilled migrants that are already in Europe, for instance skilled refugees, labor market restrictions should be lifted to ease labor market integration. If these people can be integrated fast into qualified positions without a loss in human capital, the innovation system would greatly benefit.

We have shown in this chapter that source countries can benefit from emigration through knowledge flowing back into the country. These benefits of knowledge flows can be maximized by facilitating research networks with emigrated inventors, for example by organising conferences in the origin countries. Furthermore, governments can design programs to actively keep the diaspora engaged and by encouraging and facilitating return migration. Return migrants bring back the newly gained knowledge and many times create their own start-ups which can foster development in the countries of origin.

While this chapter establishes that knowledge flows mitigate the negative consequences of emigration, further research is needed to shed light on the precise way these knowledge flows are created and characterized. Do migrants possess tacit knowledge that flows with people but cannot be transferred by other means? Or do migrants enlarge the R&D network and create better awareness of technologies in other countries? Do migrants have a competitive advantage in negotiating licensing fees with their country of origin? These open questions may guide further research in order to better understand how to increase knowledge flows and maximise their benefits.

Chapter 2

Small Steps or Giant Leaps? Competition and the Size of Innovations

2.1 Introduction

"We wanted flying cars, instead we got 140 characters."

— from Founders Fund's manifesto

This statement, often attributed to the co-founder of PayPal and Founders Fund, Peter Thiel, succinctly expresses the disappointment felt by many when comparing past expectations to the actual technological progress of the last decades. The enthusiasm after breakthroughs such as the Moon landing seems to have given way to an aspiration of quick profits through tiny improvements, for instance in communication technology with social media websites and apps. Another PayPal co-founder, Elon Musk, even felt compelled by the public's apparently waning interest in attempting "giant leaps" to other planets to pour his fortune into a risky space start-up.[1]

But is there actually any empirical evidence for this sentiment beyond anecdotes, i.e., has corporate innovation really changed in recent decades? R&D spending has increased over this time period and appears to have paid off, at least when measured by the number of

[1] The title of this chapter is a reference to Neil Armstrong's first words from the lunar surface on 21 July 1969: "That's one small step for (a) man, one giant leap for mankind." `https://www.hq.nasa.gov/office/pao/History/alsj/a11/a11.step.html` (last accessed on 5 September 2017).

granted patents. At the same time, there has been a puzzling decrease in groundbreaking innovations from corporate R&D. Arora et al. (2015) come up with a number of potential explanations for their observation that firms' innovation has indeed shifted towards more incremental innovation. One of them attributes the trend to an increase in competition due to globalization. However, the authors emphasize that their evidence only suggests a relationship, but not necessarily causality.

The present study makes three main contributions. First, the choice between innovation technologies targeting small or large inventive steps is incorporated into an R&D model with interesting implications. Second, a new measure is developed to quantify the types of innovations from patent abstracts. Third, the measure is used in instrumental variable regressions to estimate the causal relationship between import competition and the type of innovation.

The model endogenizes the allocation of resources to projects aiming at large inventive steps versus more incremental ones. Motivated by a product market competition model based on Aghion et al. (2001), I formalize how competition affects profits and, as a result, R&D portfolio choice. Competition is modeled as the extent of substitutability between the duopolists' products, but allows for a wide range of interpretations, such as trade tariffs or technological standardization. In a dynamic framework I show that an increase in the intensity of competition leads to a decline in the share of large innovations in total innovations. This effect is compatible with an inverted-U relationship between competition and the total amount of innovation. However, the negative relationship between competition and the share of large innovations holds independently of the initial level of competition. Furthermore, the intensity of competition maximizing the amount of large innovation is lower than the intensity maximizing overall innovation (which, in turn, is lower than the one maximizing the amount of small innovation). If the social value of large innovations is higher, the socially optimal level of competition is lower than the one maximizing the total number of innovations.

A firm's incentives to invest in R&D stem from two sources: On the one hand, successful cost reduction allows the firm to appropriate part of the social benefit of the innovation. On the other hand, there is an additional component to higher profits resulting from R&D that is due to business stealing (Mankiw and Whinston, 1986). In the model, market size is assumed to be fixed. Changes in competition affect business stealing more for incremental steps, where a product is very comparable to its predecessor and to the products of competitors without innovations. By contrast, for more radical innovations, this effect is negligible and the innovator effectively becomes a monopolist.[2] This study attempts to not only illuminate a crucial aspect in the relationship between competition and innovation, but also to reconcile seemingly contradictory observations such as the trend towards both more patents and fewer groundbreaking innovations. This puzzle may thus exist not (only) because of an increased ease of patenting even minor innovations over time, but because firms have adjusted their R&D strategies to increased foreign competition and produce more incremental R&D, which is less risky and leads to more, albeit smaller, steps.

The predicted relationship between competition and the type of innovation is tested using changes in competitive pressure due to growing imports from China with the data and the initial conditions instrument from Bloom et al. (2016). To quantify the type of innovation in a way that is coherent with the model, I calculate novel measures based on the similarity of a firm's patents to a comparison group of patents. The idea is that larger inventive steps lead a firm to areas in technology space which are more distant from previous innovations. Different technologies entail changes in the terms occurring in patent abstracts. Thus the continuous similarity variables may be interpreted as inverse measures of the size of inventive steps.

The estimations support the causal relationship between increasing competition and a shift towards more incremental research, i.e., the patents become more similar to a firm's own past research. While this effect is a change of firms' behavior, I also investigate whether increasing import competition affected the selection of firms which are still observed patenting in later periods, as the direction of research may be hard to adjust. The analysis of outcomes for firms with different initial R&D types suggests that the primary channel is a change in firms' behavior. There is some evidence that the sales of firms with more incremental initial research decreased less, following an increase in Chinese import competition, than those whose innovations were less incremental. Thus competition makes small steps relatively more valuable, a fact that is in line with the model.

[2] Drastic innovation implies complete business stealing, as other firms are forced out of the market by a large innovation. However, the important difference is that post-innovation profits are independent of competition. Profits after a large innovation are not affected by policies favoring competitors which can no longer be active in the market due to their technological lag.

There exists a well developed literature on the relationship between competition and innovation. Schumpeter (1942) introduces the idea that market power is crucial to innovative activity, both in terms of large firms' capability to innovate – he describes advantages to scale such as the ability to leverage the abilities of "better brains" or improved access to finance – and the incentives for "creative destruction" arising from the prospect of profits, which could not exist under perfect competition. A seminal contribution by Arrow (1962) points out the "replacement effect" in the opposite direction: A monopolist has little incentive to invent a new or improved product if this leads to self-cannibalization. I.e., it steals the old product's business, while not increasing profits as much as for an entrant, from whose perspective the entire market share of the new product is part of the innovation's net payoff. Schmidt's (1997) model shows that separation of ownership and control can give rise to another positive effect of competition on the incentives for managers through the increased threat of liquidation. However, reduced profits from a cost reduction through increased competition may lead to a non-monotonic relationship.

An inverted-U relationship between competition and innovation is also supported empirically in Aghion et al.'s (2005) influential paper. In their theoretical model, the authors find that incentives for innovation through escape competition and the Schumpeterian effect dominate at opposite ends of the spectrum of competition intensities due to the composition of industry structures (which, in their model, means duopolists may be at the same technological level or one may be ahead of the other).

Nonetheless, the survey by Schmutzler (2010) on the general relationship between competition and innovation concludes that it remains an "unresolved question".[3] Most studies in this literature do not distinguish between different types of innovations. However, a small number of publications specifically address this question and a number of further papers deal with the size of the targeted innovations in extensions.[4] The most relevant ones in the context of this study are briefly summarized below.

[3] An earlier survey by Gilbert (2006) similarly describes the conclusions from the literature as "meager".

[4] Note that this is different from the distinction between process and product innovation. Either of these innovations may be small or large, as the new product or process may be more or less similar to what has existed before. However, in my model innovations affect only the current market through cost reductions, such that process innovation is the appropriate concept. Nonetheless, in reality the distinction is less sharp and a new product, such as the smartphone, may render products in various other markets obsolete (such as MP3 players and compact digital cameras).

Cabral (1994) is the first paper to study this question and shows that firms in the market will take on less risk in their R&D portfolios than would be socially optimal due to the effect of the competitor's success probability on the difference between social and private returns. Aghion and Howitt (1992) come to a similar conclusion in an endogenous growth model, in which the result is driven by the effect of business stealing. However, none of these papers focuses on the *change* in this bias when the intensity of competition increases. In a later paper, Cabral (2003) studies the effect of a firm's position on the quality ladder on its R&D strategy. He finds that even with an infinite horizon firms will play the risky strategy when sufficiently behind and the safe one when they lead. The critical assumption is that there is a lower and an upper limit to a firm's profit, such that there is indeed "less to lose" when behind than when ahead. There is no intensity of competition modeled beyond the technological position, which does not capture a parameter that competition or trade policy can directly affect. Kwon (2010) studies the influence of the number of competitors on firms' R&D choice between a risky and safe project. Assuming that each of the two innovations can be patented only by a single firm, which is drawn at random among the innovators, Kwon finds a bias towards low risk projects as well. He also analyzes the change of this bias with an increase in competition, which in his model means a higher number of competitors. Understanding competition in this way, the study predicts that more competition makes riskier projects more attractive due to the lower probability of parallel innovation. In my model, by contrast, the number of firms is fixed and competition only affects their profits.

A distinction between exploration and exploitation has also been made more recently by Manso (2011) and a series of following papers which further explore the topic (e.g. Ederer, 2013; Ederer and Manso, 2013; Balsmeier et al., 2017). These studies approach the question from an organizational perspective, abstracting from the market structure that determines which type of innovation the owner of a firm wishes to incentivize for an agent or a team to choose in the first place. Their models feature a principal-agent framework, in which an optimal incentive contract is determined, whereas my study is concerned with the effect of the incentives that emerge in the market on different types of R&D investments. Connecting these two elements and exploring the transmission of market incentives to contracts and organizational incentive structures is an interesting area for future research, especially since R&D managers require psychological elements to be taken into account.[5]

While the interpretation of the competition parameter in this and similar models is rather abstract, the empirical part of the chapter focuses on one particular type of competition that has received increasing attention recently: the competition that comes with openness to trade and in particular the effect of Chinese import competition. Various recent articles have studied the impact of this "China Shock" (Autor et al., 2016b) on American and European economies. This surge in imports from China has been the result of increasing openness to global markets by the country's government, in particular through broad reforms in 1993 to enable the country to become a WTO member (Naughton, 2007). The accession followed in 2001.

Whereas much of this literature has been concerned with the effects of these imports on the domestic labor markets (Autor et al., 2013; Acemoglu et al., 2015) and a diverse set of other questions, two recent studies in this literature, which explore the effect on innovation, come to opposite conclusions: Bloom et al.'s (2016) study finds that European firms are driven towards increased technological upgrading and more innovation by competition (an effect "within firms"), as well as an increased chance of survival for firms which are already technologically ahead. The empirical part of my study builds on their work, combining their data with additional data from PATSTAT, the Worldwide Patent Statistical Database of the European Patent Office. Autor et al.'s (2016a) later working paper finds that innovation has decreased in American firms as a result of increased imports from China. While several reasons for the opposite sign of the effect compared to Bloom et al.'s (2016) work are discussed, in reference to Aghion et al. (2005) the authors suggest that the more competitive markets in the U.S. may simply be beyond the maximum of the inverted-U, such that additional competition decreases innovation.

The next section of the chapter presents the model and derives its implications for the effects of competition on the types of corporate innovation. Section 2.3 describes the empirical

[5] Lacetera and Zirulia (2012) combine both approaches, however the choice between basic and applied research in their model does not reflect the feature of interest to my study. In their paper, the distinction between the two types is that basic research is open, which leads to externalities and from which the researcher derives utility, while in my study the larger inventive step of more radical (and possibly more basic) research endeavors is essential.

approach and the construction of the measures for the types of innovation, as well as the empirical results. Section 2.4 summarizes the conclusions and outlines potential directions for future research. The appendix presents a model without the simplifying assumptions that enabled an analytical solution of the main model. Furthermore, additional empirical robustness checks are shown.

2.2 Model

2.2.1 Set-Up

Incentives for different types of R&D investments depend on the difference between product market profits with and without the innovations. This section first presents the set-up of the dynamic R&D model and the assumptions regarding profit levels with different technologies and their changes with competition. In particular, while being ahead by a large step always yields higher profits than by a small step, it is assumed that higher competition increases profits only for a firm leading by a small step. A large step yields constant monopoly profits. The motivation for these features is based on models from Aghion et al. (2001, 2002)[6] and is discussed after the presentation of the model. These assumptions are then shown to imply a decreasing share of large innovations in all innovations. This decreasing share is compatible with Aghion et al.'s (2005) inverted-U relationship between competition and total innovation and holds independently of the initial level of competition, i.e., the position on the curve. The maximum of the absolute level of large innovation then occurs at a lower level of competition than that of total innovation. Thus the model suggests that maximizing total innovation may not be the socially optimal policy if large steps lead to higher externalities.

In the R&D model, there are two firms competing in an industry over an infinite time horizon. In each period, innovations of different types and firms are disjoint events, such that at most one may occur, and the probabilities of small and large innovations for the two firms depend on their investments in the previous round. The R&D outcome determines the relative

[6]Note that only the working paper of Aghion et al.'s (2002) study models the product market explicitly. The published version (Aghion et al., 2005) abstracts away from the details of product market competition and instead introduces a parameter measuring the extent to which the duopolists are able to collude.

technological positions of the firms and their product market profits. Then R&D investments are made and the next period follows. The firms are assumed to be myopic in the sense that they consider only next period's product market profit and current R&D costs in their optimization.

The industry state is characterized by their relative positions: They may be at the same technological level or one firm may be ahead of the other, either by a small or by a large step. Following Aghion et al. (2005), it is assumed that the leading firm cannot invest in R&D. A firm leading by a large step earns a profit of $\pi_L = 1$. Irrespective of the size of the lag, the follower can invest only in catch-up R&D and earns a product market profit of $\pi_F = 0$. The intensity of competition is captured by the variable $\Delta \in [0, \frac{1}{2}]$ and affects product market profits for leveled firms and small step leaders. Each leveled firm earns $\pi_0 = \frac{1}{2} - \Delta$, i.e., higher competition reduces the sum of their profits from 1, which equals a large step leader's profit, to 0. A firm leading by a small step earns $\pi_S = \frac{1}{2} + \Delta$. Thus the firms' profits increase from $\frac{1}{2}$ to 1 as competition increases. Only firms in leveled industries have two continuous choice variables, n_L for large steps and n_S for small steps. The variables n_i, $i \in \{L, S, F\}$ are the success probabilities for each type of innovation in a period and are associated with costs[7] $2n_i^2$, $i \in \{F, S, L\}$. Since it is assumed that there are no parallel innovations, R&D costs and profits ensure that the sum of probabilities in equilibrium does not exceed 1. In addition, the laggard's benefits from spillovers are reflected in the help parameter h (with $h < \frac{7}{8}$).[8] In terms of model dynamics, this parameter allows an exogenous influence on the tendency of the economy to return to the leveled state (or, equivalently, the probability that an industry is in the leveled state at any point in time). The help parameter can be seen as an additional policy variable reflecting the strength of intellectual property protection and the extent of

[7]Since this study is concerned primarily with the *changes* in investments in the two innovation technologies rather than their levels, in the model the equilibrium probability of achieving a large innovation will never be smaller than that of a small innovation. This is the case simply because cost functions are assumed to be the same for both technologies and thus a large cost reduction is more attractive by definition. The appendix shows that the levels may be easily adjusted by introducing a cost parameter, such that large innovations are more costly. Since the mechanism is not affected by this, I opted for improved analytical tractability in the main part of the chapter and omitted these cost parameters.

[8]This help parameter benefits the laggard by increasing the probability of catching up and earning the profit π_0 in the leveled state next period. Such assumptions on h are not necessary in the model in the appendix, since the probability of two innovations occurring at the same time is 0 in continuous time. The model in the appendix supports the results of this section.

Figure 2.1: Possible Positions on the Technology Ladder

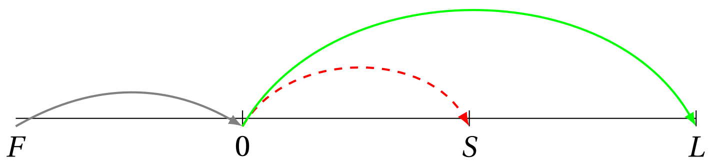

Notes: Firms may be equally advanced (0), lead by a small step (S) or by a large step (L), or lag behind as follower (F).

knowledge spillovers between companies.[9] These positions[10] and the possible movements between them are depicted in Figure 2.1.

Since the model abstracts away from strategic interaction and firms look ahead only one period, a firm's profit and R&D investment depend only on its technological position and the intensity of competition. The maximization problem for firms in the leveled state is

$$\begin{aligned} V_0 = \max_{\{n_L, n_S\}} \ & \pi_0 - 2n_L^2 - 2n_S^2 \\ & + n_L\pi_L + n_S\pi_S + (\overline{n}_L + \overline{n}_S)\pi_F + (1 - n_L - n_S - \overline{n}_L - \overline{n}_S)\pi_0, \end{aligned} \tag{2.1}$$

where π_0 is the product market profit of a firm in the neck-and-neck state, from which the total costs of R&D are subtracted. After the current period, the firm may find itself in a different state. For simplicity, the next period is not discounted and its value is the expected profit in the product market. Furthermore, as it is assumed that only either one or no innovation may occur in each sector per period, parallel innovations are not considered. With probability n_L the firm makes a large innovation and with n_S a small one, leading to profits π_L and π_S, respectively, in the product market. The next term takes the potential transition to position F into account: If the other firm innovates, whose investment variables are $\overline{n}_L$ and $\overline{n}_S$ and are taken as given by the optimizing firm, the future profit will be that of a follower π_F. Finally,

[9]An example of a policy likely to affect these spillovers are non-compete contracts, which have been found to even affect relocation decisions of knowledge workers by Marx et al. (2015). As a result of the assumption that firms look ahead only one period, increasing h always increases innovation in this chapter's model. If firms take all future periods into account as in the model in the appendix, however, faster catch-up decreases the value of being ahead.

[10]From the perspective of the following firm, the size of the lag does not matter such that only one position is depicted.

none of these transitions may occur and the next period's present value will be the same neck-and-neck present value as in the current period.

A lagging firm only has to choose the level of catch-up R&D n_F in its optimization (where n_F is the probability of catching up, for simplicity irrespective of the lag size), such that its objective function is:

$$V_F = \max_{n_F} \pi_F - 2n_F^2 + (n_F + h)\pi_0 + (1 - n_F - h)\pi_F \tag{2.2}$$

The structure of the objective function V_F is similar to that of leveled firms, except that there are fewer choices to make and the help parameter h plays a role in influencing the probability of a successful catch-up. Finally, a firm leading by a small or large step has no decisions to take and no R&D costs, such that there is no R&D optimization.

Firms' maximization leads to the first order conditions below, which give the optimal investment levels.

$$4n_L = \pi_L - \pi_0 \tag{2.3}$$

$$4n_S = \pi_S - \pi_0 \tag{2.4}$$

$$4n_F = \pi_0 - \pi_F \tag{2.5}$$

To be consistent with a steady state, the probabilities that an industry is in either of the three states (leveled with probability μ_0 or unleveled with probability μ_F, where one firm is leading by either a large step or by a small step) have to sum to one: $\mu_0 + \mu_F = 1$. Furthermore, the inflows and outflows of the states have to be equal.

$$2\mu_0(n_L + n_S) = \mu_F(n_F + h) \tag{2.6}$$

The average innovation rates I_L for large and I_S for small steps in the economy depend on the investment choices of firms in the leveled state, as well as the probability of the industry being in that state μ_0. The decrease of this latter probability for high levels of competition is the source of the downward sloping part of the inverted-U relationship between competition and innovation.

$$I_L = 2\mu_0 n_L \tag{2.7}$$

$$I_S = 2\mu_0 n_S \tag{2.8}$$

These innovation rates describe the expected number of innovations per period for the average industry for large and small steps, respectively.[11] The results below describe how these rates are affected by the intensity of competition Δ.

2.2.2 Discussion of the Model

In order to focus most clearly on the mechanisms leading to the results, the model makes several assumptions. The most important ones are discussed in this section.

The model summarizes the product market in the profit levels and, in particular, assumes that leveled firms' profits decrease in competition and a small step leader's profits increase, while profits after a large step are constant. These features are consistent with models such as Aghion et al. (2001), as explored further in Appendix B.4. These models assume that consumer spending is fixed in a duopolistic industry. Competition is modeled as the substitutability between the firms' products. The size of the inventive step is reflected in the cost ratio. In the extreme case where a large innovation leads to zero costs, the innovator serves the entire market and its profit equals consumer spending, which is 1, independently of competition. The idea is that if a technological advance is sufficiently large, the other firm will leave the market and thus policies, which make it easier for a competitor to steal business, will have no effect. Increasing competition decreases the profits of leveled firms, however, and it

[11] The innovation rates could also be interpreted as the probability of an innovation of a certain type in a period. E.g. the probability of a large innovation (I_L) depends on the probability that the industry is in the leveled state (μ_0), because this is necessary for innovations at the technological frontier. Furthermore, I_L depends on the probability of producing a large innovation in the leveled state for each firm (n_L). This is multiplied by two, because there are two firms and innovations are disjoint events.

increases the benefits of leading by a small step.[12] A small technological improvement by one competitor is not sufficient to take over the entire market, e.g. because of barriers such as tariffs, transportation costs or incompatibility. Thus reducing such barriers is assumed to have a bigger effect on competitors which are technologically close. The empirical analysis below provides suggestive evidence that small steps indeed become more beneficial under higher competition.

To focus on the changes in industry structure and innovation, the model assumes that there are no parallel innovations. This simplifies firms' optimization and could be justified by the fact that, as the period under consideration becomes shorter, success probabilities decrease and the probability of parallel innovation may become negligible.[13] The assumption that firms look ahead only one period can be explained by the limited tenure of CEOs and the insufficient alignment of their incentives with long-term investors' interests.[14] The considerably more complex model in the appendix does not make this assumption and the results survive in the numerical solution. An assumption that remains, however, is that while firms maximize their present value through their R&D investments, they are unable to commit to future implementation of strategies that are not profit maximizing from the perspective of another position, in which the firm might find itself later.[15] The exclusion of such strategic considerations can again be justified by a high probability that CEOs will be replaced if they do not maximize the company's present value at any point in time. Thus the behavior of the future company, whether decided by the same CEO or another, is constrained and has to be taken as given by the current CEO.

The fact that the number of firms in the model is fixed may be justified by the existence of strong barriers to entry in many innovative sectors. In other markets, however, the effect on entrepreneurs' innovation should factor into competition policy. While the optimal level of

[12]The model in the appendix assumes that small steps are sufficiently large for the leader's profits to increase with competition. Besides the arguments discussed in the context of the model in the appendix, it should be noted that this model's emphasis is on comparing drastic to all other innovations. The exclusion of smaller advances also seems justified in view of the empirical analysis: Working with patent data, a truncation of observed inventive steps is inevitable, i.e., the smallest improvements will not satisfy patentability requirements.

[13]The two innovation technologies for the two firms could also be thought of as leading into four different technological directions and only at most one is "correct" and leads to an improvement.

[14]This simplification is also made in Chapter 12 of Aghion and Howitt (2008).

[15]A leader, for example, might want to threaten higher R&D investments in the leveled state, than is profitable at that position, in order to deter the other firm from investing in catch-up R&D.

competition may change, the shift towards small innovations should persist if outside innovators can replace incumbents and compare post-innovation profits. Competition increases the profits which entrants can expect after a small innovation, while large innovations yield constant monopoly profits. Similarly, if outsiders need to enter through catch-up innovation, their incentives upon entry would be shifted towards small steps by competition.

2.2.3 Results

The innovation rates I_L and I_S, as defined above, take only innovations at the technological frontier into account, while catch-up innovation does not push beyond the best available technology in the industry. R&D investments in these types of innovation are therefore made only when firms are at the same technological position. Their success probabilities for both types of innovation (n_S for small steps and n_L for large steps) are shown in Figure 2.2. Both investments increase in competition (Δ) because pre-innovation profits (i.e., profits in the leveled state) decrease and incentives are determined by the difference between post-innovation and pre-innovation profits. The difference in the slope is the result of increasing profits for the small step leader, but constant profits after large innovations. Figure 2.2 also shows that in the unleveled states the follower's investment in catching up (n_F) decreases with competition, as reaching the leveled state becomes less attractive.

Proposition 1. *The share of large steps in all innovations in the economy decreases with the intensity of competition.*

$$\frac{d}{d\Delta}\left(\frac{I_L}{I_L + I_S}\right) < 0 \tag{2.9}$$

Proof. See appendix, Section B.1. □

The proposition states that the share of large innovations declines for any level of competition. The change in this share with competition can be seen in Figure 2.3. The intuition is that more competition makes both types of innovation more attractive as it decreases a neck-and-neck firm's profits. However, only small innovation additionally becomes more attractive because of increased post-innovation profits through additional business stealing. The inverted-U relationship between total innovation and competition is maintained in the numerically solved model in the appendix, in which profits follow directly from a product market model.

Figure 2.2: R&D Investments in Leveled State

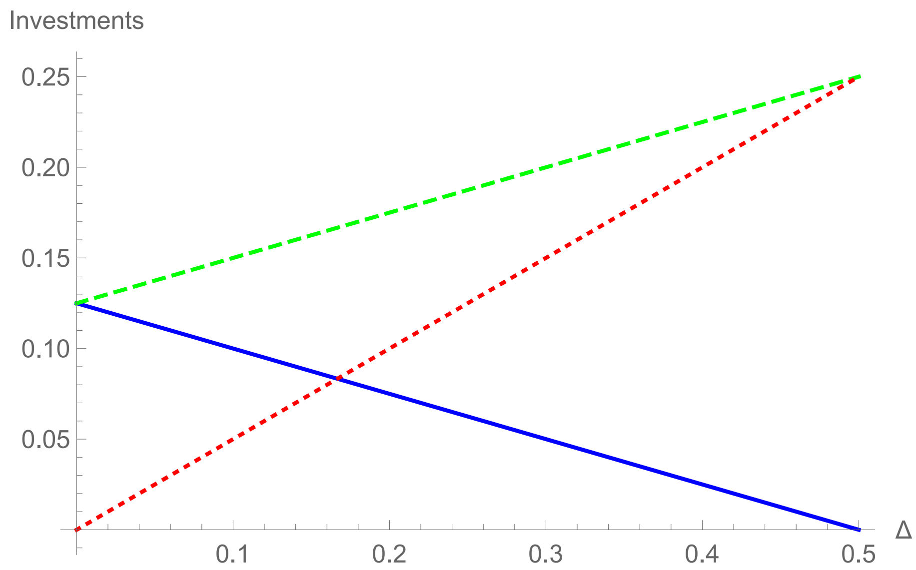

Notes: Investments in R&D of each type for a firm in the leveled industry state (n_L in dashed green and n_S in dotted red) and in catch-up innovation for unleveled states (n_F in solid blue). ($h = 0.2$)

Figure 2.3: Share of Large Innovations

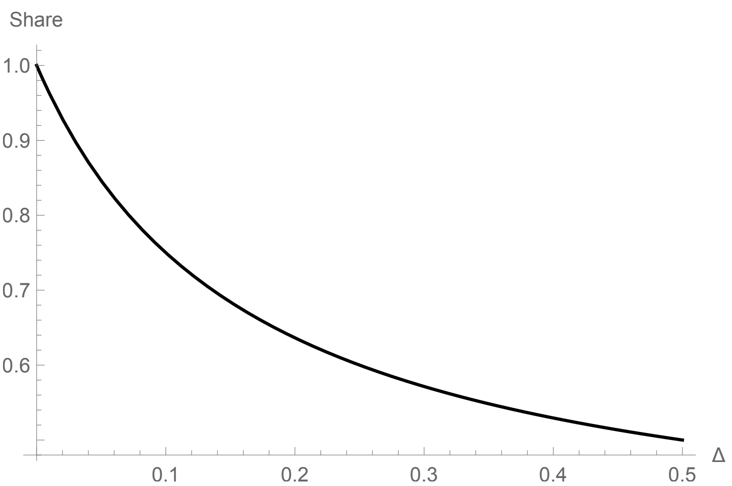

Notes: This graph shows the share of large innovations in all innovations $\frac{I_L}{I_L+I_S}$ depending on the intensity of competition Δ. ($h = 0.2$.)

Figure 2.4: Absolute Levels of Innovation

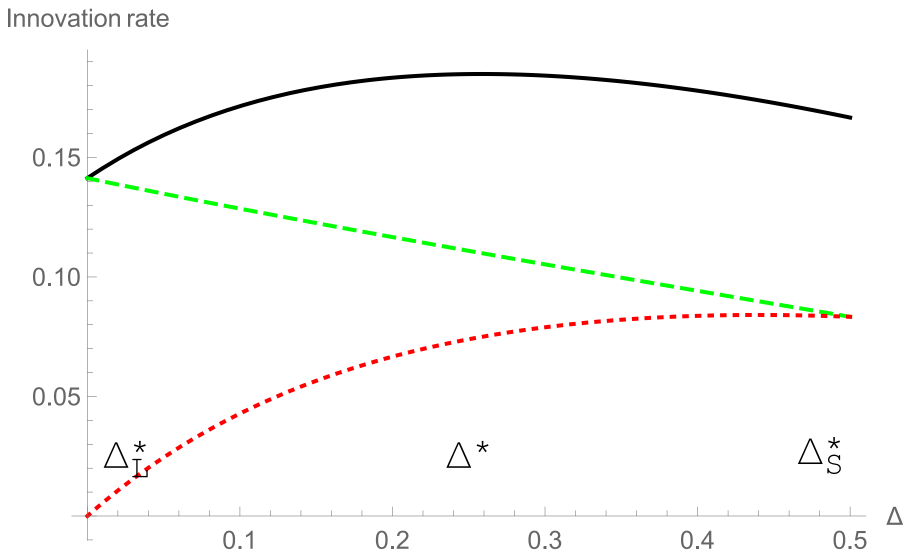

Notes: This graph shows the absolute levels of total innovations I, as well as large (I_L) and small (I_S) innovations, for a given intensity of competition (Δ): I in solid black, I_L in dashed green and I_S in dotted red. ($h = 0.2$)

As Figure 2.4 demonstrates, for some parameter values, the amounts of small and large innovation (i.e., innovation rates I_S and I_L) may themselves follow an inverted-U, not only their sum. (In the figure, this is only the case for I_S.) The fact that the inverted-U relationship between overall innovation and competition can persist indicates that there is no conflict between this study and the results of Aghion et al. (2005). Rather, this model investigates the composition of the total amount of innovation, which is more commonly studied. Nonetheless, as described below, while the socially optimal level of competition may still be interior (i.e., $\Delta \in (0, \frac{1}{2})$), it need not be the level that maximizes overall competition, when externalities of different types of innovation differ. Note that unlike the effect of competition on the amount of innovation, the effect on the *share* of large innovations stated in Proposition 1 is negative, independently of the initial level of competition in the model. Thus, while the effect of an increase in competition, e.g. due to globalization, on total innovation may be positive or negative, the model unambiguously predicts a shift towards smaller steps.[16]

[16] A statement from Google executives Schmidt and Rosenberg (2014) fits well with the model's prediction that, as globalization has increased competition, the share of resources devoted by firms to catch-up and small step innovation has increased. In their book "How Google Works" they lament: "Business leaders spend much of their time watching and copying the competition, and when they do finally break away and try something new, they are careful risk-takers, developing only incremental, low-impact changes."

Aghion et al. (2005) explain the inverted-U relationship between competition and innovation as the interplay of two countervailing effects: On the one hand, more competition increases the difference between pre-innovation and post-innovation profits in the leveled state ("escape competition"). On the other hand, the incentives for a laggard to catch up to the technological frontier decrease as competition becomes more intense ("Schumpeterian competition"), as is easily seen in the extreme case of Bertrand competition, where both, lagging behind and being at the same level of the competitor, yield zero profits. They assume that a leader cannot benefit from additional innovation, e.g. because of spillovers, such that it never invests in R&D. A crucial feature of their model is the "composition effect", which results from the decreasing share of industries being in the leveled state, rather than in the unleveled state, as competition increases, and which determines which of the two previous effects dominates. For low levels of competition, an industry is more likely to be in the leveled state, such that the positive escape competition effect on neck-and-neck firms' innovation prevails. For high levels of competition, an industry is more often in the unleveled state and further increases in competition make catching up and reaching the leveled state less attractive for the follower (while the leader does not innovate at all).

Proposition 2. *The absolute level of small innovation I_S follows an inverted-U shape for $h < \frac{1}{4}$ and otherwise is increasing in competition for any initial level of competition. The absolute level of large innovation I_L is decreasing in competition if $h \leq \frac{1+\sqrt{6}}{8}$ and increasing if $h \geq \frac{1+\sqrt{5}}{4}$. If h is in between these values, I_L has an inverted-U shape.*

Total innovation $I = I_L + I_S$ is increasing in competition if $h \geq \sqrt{\frac{1}{6}}$ and is otherwise inverted-U shaped.

Proof. See appendix, Section B.2. □

Intuitively, the reason for the inverted-U is that incentives to innovate in the leveled state increase with competition, but the share of industries able to innovate at the technology frontier decreases as fewer industries are in that state. Figure 2.5 shows the share of leveled industries μ_0 as a function of competition.[17]

[17] The parameter h helps the follower to catch up. Hence, the higher this help parameter is, the more likely is it that an increase in competition increases innovation, since it increases the probability of being in the leveled state, in which competition increases R&D investments.

Figure 2.5: Share of Industries in Leveled State

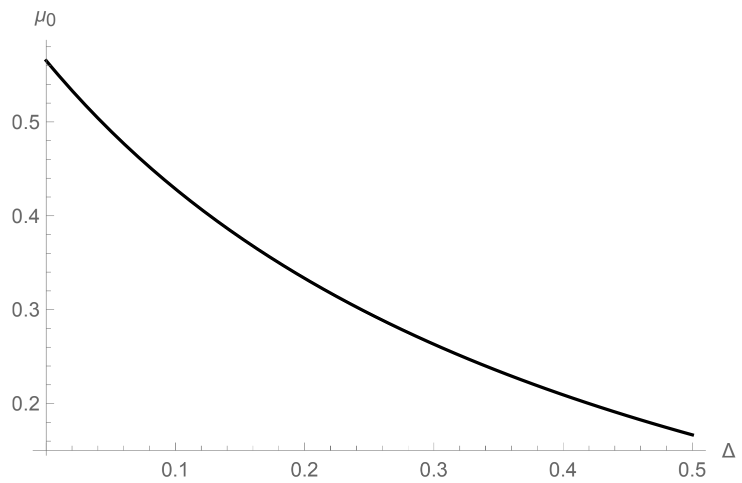

Notes: Share of industries in the leveled state for a given intensity of competition (Δ). ($h = 0.2$)

While for a model with one type of innovation the level of competition at which the curve reaches its maximum may be of interest for the optimal policy, a question that arises only in this chapter's model is the relative position of the curves' maxima. In Figure 2.4 total innovation is maximized at the interior competition intensity Δ^*, while the expected number of small innovations per period I_S is maximized at the higher level Δ_S^*. By contrast, the innovation rate for large innovations is maximized at the lowest level of competition ($\Delta_L^* = 0$).

Proposition 3. *The level of competition Δ_L^*, at which the absolute amount of large step innovation the economy produces is maximized, is at most Δ^*, which maximizes total innovation. The level Δ_S^*, which maximizes the amount of small step innovation, is at least as high as Δ^*.*

$$\Delta_L^* \leq \Delta^* \leq \Delta_S^*$$

Proof. See appendix, Section B.3. □

The proposition confirms that this order of the competition intensities maximizing large, total and small innovation holds generally. The variable of interest to policy makers – the socially optimal level of competition – thus depends on the relative social value of small and large innovations. Taking a step back from this abstract model, this means that the fact that small and large innovations are likely to have different net social values, the optimal policy differs from an unweighted maximization of the number of innovations, e.g. approximated by the number of patents. If larger innovations tend to have higher externalities, the optimal level of competition could be lower than the one that maximizes overall innovations. The optimal competition policy ideally takes technological differences into account. Established industries, in which the social benefits of further large step innovations seem low, may be regulated more tightly, whereas lenient regulation may be appropriate for nascent industries with new technologies, where experimentation is relatively more likely to yield sizable positive externalities on consumers and other innovators.

Proposition 2 shows that the model is compatible with the inverted-U relationship found in the literature. Proposition 3 demonstrates that the effects of competition on different types of innovation are relevant for the optimal policy. The main prediction of the model, however, is Proposition 1.

Hypothesis: An increase in the intensity of competition decreases the average size of corporate innovations.

This hypothesis is tested in the next section. The results from the model are driven by the effect of competition on profits at different technological positions.[18] Below, the impact of the type of inital innovation on various economic outcomes for firms is tested. The empirical results suggest that competition indeed increases the relative value of small steps through higher sales.

2.3 Empirical Identification of the Relationship between Competition and the Type of R&D

The model presented in the previous section and the research question itself contain variables which pose an empirical challenge not only because of identification, but also because of measurement. This section discusses the study's approach to quantifying competition and,

[18] For simplicity, the relationship between competition and product market profits with different relative technological positions is an assumption in this section, which then explores the implications. By contrast, in Appendix B.4, a similar relationship emerges from a product market competition model based on Aghion et al. (2001). The empirical results are also in line with this model. Higher competition enables the leading firm to sell more with given marginal costs.

more specific to this context, the type of innovation, so as to avoid "making assumptions about how unmeasurable things affected other unmeasurable things".[19]

There are various ways to measure competition and the substitutability parameter α in models like Aghion et al. (2001) may be interpreted as representing different policies affecting competition. One important area in which policy makers can influence the extent to which firms are able to enter each others' market is trade policy. For instance, Autor et al. (2016b) suggest the inverted-U relationship between competition and innovation from Aghion et al. (2005) as a potential explanation for the contradiction between their results and Bloom et al.'s (2016), thereby interpreting competition in the R&D model as import competition. While production processes may well differ between Chinese exporters and their European competitors, e.g. due to lower labor costs, the matching industry classification in the data set suggests that they serve the same markets, such as the one for "Vehicular Lighting Equipment" (SIC 3647). Note that the analysis includes only patents of European firms and that their market position may be threatened even if the imported goods are of lower quality, as long as they are sufficiently cheaper. The model's intuition is still reflected in this setting, since a small improvement in quality (or a decrease in production costs) could then be more valuable by preventing business stealing. A sufficiently large technological advantage always allows a firm to sell to the entire market, but the increased availability of imported goods in the market makes a small step relatively more likely to attract customers.

A bigger challenge is posed by the type of innovation, which far fewer empirical studies have tried to quantify. To get a continuous measure of a patent's similarity compared to previous research I use natural language processing to evaluate patent abstracts. Such machine learning methods are gaining popularity in economics in general, as well as in innovation economics in particular[20]. Hoberg and Phillips (2010) and Kaplan and Vakili (2015) may be the methodologically most closely related papers. The former study uses the similarity of product descriptions to study mergers and acquisitions and the latter constructs a binary

[19] In one of Paul Krugman's New York Times columns from August 2013, he argued that a reason for why New Growth Theory fell out of favor was that "too much of it involved making assumptions about how unmeasurable things affected other unmeasurable things". URL: https://krugman.blogs.nytimes.com/2013/08/18/the-new-growth-fizzle/ (last accessed on 30 May 2017).

[20] Recent examples of textual analysis methods and applications in economics include Baker et al. (2016), Hansen et al. (2017) and Hansen and McMahon (2016). For overviews of available methods see Balsmeier et al. (2016) and Bholat et al. (2015).

measure of novelty, which is calculated using a latent Dirichlet allocation (LDA) (Blei et al., 2003), to evaluate the "cognitive novelty" of innovations. My preferred measure, which is described below, uses a different approach, but in the appendix I verify that the results are qualitatively unchanged when using LDA, among other variations in the calculation of the similarity measure. Kaplan and Vakili (2015) is not only relevant to this study because of their technical approach, but also because of the economic insight gained from their analysis: Their study draws a clear distinction between the economic value and the "cognitive novelty" of an innovation, which may be viewed as corresponding to profitability in the product market and step size (or cost reduction in the model in the appendix) of an innovation, respectively. Hence, I estimate how the average novelty in a market is affected by import competition. The results suggest that this choice is indeed driven by a change in the private economic value of novelty.

Another advantage of the similarity measure over citation-based measures is that it is fixed at the time of the application, whereas forward citations accumulate afterwards as further innovations build on the patent. The extent of follow-on innovation, however, may itself be the result of the economic value rather than of cognitive novelty.[21] While the two are clearly correlated, only the value, but not the novelty, of a given innovation can be affected by competition ex post.

2.3.1 Data and Measure Based on Latent Semantic Analysis of Patents

This study presents a first empirical test of the causal relationship between product market competition (more precisely, Chinese import competition) and the *type* of innovation firms produce. Each firm in the sample is assigned a novel measure of R&D type (interpreted as the extent to which the R&D portfolio is geared towards incremental innovation rather than more radical or drastic innovation) based on textual analysis of the abstracts from its patent filings in a year. This section explains the steps to construct the measure of the type of corporate innovations used in the empirical analysis below from patent data from the European Patent Office (PATSTAT 2014 Autumn Edition, which includes not only data on European patents but also from various national patent offices and in particular the USPTO).

[21] In fact, Kaplan and Vakili (2015) use forward citations as their measure of economic value.

To prepare the corpus of documents (i.e., patent abstracts) for further analysis, I limit the sample to abstracts in English[22] and use the Natural Language Toolkit (Bird et al., 2009) to stem all words with the Porter Stemmer (Porter, 1980), such that ideally inflections of the same word result in the same stem. In a next step, the term frequency-inverse document frequency (tf-idf) matrix is produced using scikit-learn, a machine learning library for the Python programming language (Pedregosa et al., 2011). This approach is commonly used in the area of information retrieval for natural language processing and transforms the text in each document into a vector, in which each word (or in this case, more precisely, word stem) is one dimension. Thus, the order of words in the abstract is ignored (a so-called "bag-of-words model"). By contrast to raw frequencies, tf-idf assigns lower weights to words that appear in a large share of documents (like "signal", "method" or "process"), whereas terms that appear in a smaller subset of abstracts (like "auction", "penicillin" or other more specific terms) are assigned a higher weight in this measure.[23] This method better approximates an emphasis on the terms that define the content of a document. Furthermore, English stop words (e.g. articles like "the") are ignored explicitly, as well as stems occurring in more than 50% of abstracts or in fewer than 5 abstracts (not only in an individual firm's patents, but in all of the patents over the entire sample period).

While the tf-idf matrix could be used already to measure the similarity between documents, e.g. by calculating the cosine similarity of vectors, I next apply a singular value decomposition (Deerwester et al., 1990).[24] SVD is applied for dimensionality reduction and used as part of latent semantic analysis (LSA, sometimes also called latent semantic indexing), as a reduced matrix with fewer dimensions is better able to capture synonymy and polysemy, i.e., words

[22] If patent applications have multiple abstracts, PATSTAT contains the English version according to its data catalogue. For a few abstracts, however, the language is incorrectly classified as English in the data. I remove these observations by verifying the language with the "Compact Language Detection in R" package, which makes Google Chrome's language identification library available for the R programming language (`https://cran.r-project.org/src/contrib/Archive/cldr/`).

[23] Implementations of tf-idf in different studies and software packages vary in their precise functional form. Here, a term's tf-idf value is $\underbrace{n_{d,v}}_{tf}\underbrace{(1+ln(\frac{1+n_d}{1+df_v}))}_{idf}$, where $n_{d,v}$ is the number of occurrences of term v in document d, n_d is the number of documents and df_v is the number of documents which contain v. Vectors are then divided by their Euclidean length for normalization. The results are robust to using sublinear scaling for term frequencies instead, which means that a term's additional importance from additional occurrences is decreasing (Manning et al., 2008) and the tf-idf expression becomes $(1+ln(n_{d,v}))(1+ln(\frac{1+n_d}{1+df_v}))$ (and tf-idf is set to zero for terms that do not occur).

[24] Table B.16 shows that the results are comparable if similarities are calculated directly from a tf-idf matrix.

with multiple meanings (Manning et al., 2008; Murphy, 2012). Intuitively, the document vectors from the tf-idf matrix, which contains 70167 documents and 17922 terms, are "summarized" to a smaller number of components in this low-rank approximation based on SVD. While it is not possible to interpret the meaning of each of the 100 continuous variables (i.e., components) that are calculated for each abstract, the main advantage of SVD for this study is that different words with related meaning tend to be mapped into the same dimensions, such that a comparison based on SVD is more robust to changes in wording without changes in meaning.

Next, the vectors with the 100 components for each patent of a firm filed in a year are averaged and their cosine similarity to the average firm's vector of the comparison group is calculated as the similarity measure. The comparison group consists of patent applications filed in the preceding four-year period[25] and includes patents of either only the same firm, of all firms in the same two, three or four-digit SIC category, or of all firms in the sample.

In short, a patent whose abstract is very similar to the comparison group's corpus of patent abstracts is assigned a high similarity score.

These similarity measures are constructed for data from Bloom et al. (2016) on European firms and their exposure to Chinese import competition, to which they are matched using the Amadeus-PATSTAT-match provided by Peruzzi et al. (2014).[26] Bloom et al.'s (2016) data sources include Bureau van Dijk's Amadeus database on firms and UN Comtrade. In the regressions I use matched patent applications from all patent offices covered in PATSTAT except for Chinese patents to maximize the number of observations, while excluding direct effects on the propensity to protect intellectual property in the exporters' own market. Another reason for the exclusion of patents filed with the Chinese patent office (SIPO) from the analysis is the sharp increase in patenting there starting in the late 1990s (Hu and Jefferson, 2008). The literature suggests that several factors have contributed to this "Chinese patent explosion" and that changes in patent law as well as lower novelty requirements may be among them (Hu and Jefferson, 2009; Eberhardt et al., 2016). Robustness checks in the appendix with all

[25]The comparison group consists of four years to increase the probability that a firm has comparison patents, such that it can be included in the regressions, while excluding the earlier year for which the similarity measure is used in the long differences estimation.

[26]I retain only observations with a high match quality (variable "phat" in Peruzzi et al.'s (2014) dataset of more than 0.99).

patents (including patents filed in China and those with abstracts in non-English languages) or selections on other criteria show that the qualitative results are not sensitive to these choices.

2.3.2 Econometric Specification and Results

The econometric strategy is based on Bloom et al. (2016), who study the effect of Chinese import competition on the number of patents, IT upgrading and total factor productivity (TFP). Combining the data as explained in the previous section, I estimate the following regression equation:

$$\Delta \text{similarity}_{ijkt} = \beta \Delta \text{ShareImChina}_{ijkt} + \Delta f_t + \varepsilon_{ijkt}, \tag{2.10}$$

where indices denote firm i, industry j, country k and year t. The equation is estimated in five-year differences, such that fixed effects that remain constant over time, e.g. for firms, industry sectors and countries, disappear. The regression includes year fixed effects Δf_t to account for general time trends which should not be attributed to differential growth in Chinese import penetration. $\Delta \text{ShareImChina}_{ijkt}$ is the regressor of interest – the five-year difference in the share of imports to an industry coming from China – and can be viewed as roughly corresponding to the extent to which firms are able to enter each other's markets and steal business through superior or cheaper products, which is captured in the competition parameter in the model.[27] While import competition data are available at the four-digit SIC level, firms' exposure depends on the distribution of their activity over (one or several) industries, such that the equation is estimated on the firm level and not pooled at the industry-level. As most of the variation is between four-digit SIC categories (rather than between firms within an industry), standard errors are clustered on the SIC4 level.

[27] Note that the Δ in the model measures competition and can be viewed as a function of substitutability α in the model in the appendix, whereas in this empirical section "Δ" denotes five-year differences in the variables.

As Chinese imports may have increased more in some sectors *because* its innovations were more or less incremental than others, I use Bloom et al.'s (2016) initial conditions instrument[28] for the endogenous variable $\Delta\text{ShareImChina}_{ijkt}$. The two-stage least squares regression reported below instruments for this variable with $\text{ShareImChina}_{jt-6} \cdot \Delta\text{ShareImChina}_{t}$, i.e., the initial share of Chinese imports of all imports in the specific sector (in the European Union and the U.S.) is interacted with the overall growth of Chinese import competition. Appendix B.5.3 reports first-stage estimation results and OLS regressions. The first stage shows that the instrument is sufficiently powerful despite the reduction in sample size compared to Bloom et al. (2016). The F statistic is above the rule of thumb of 10 for weak instruments (Staiger and Stock, 1997). Amiti and Freund (2010) find that Chinese imports increased mostly in existing products rather than new ones. Bloom et al. (2016) argue that this supports the instrument's exogeneity. This also explains the instrument's relevance and, in the context of my study, suggests that import growth was not driven by Chinese exporters' expectations about the development of particular industries' innovation. Note that the exclusion restriction relies on the *changes* in similarity not being directly affected by initial conditions (or both being affected by a third variable), while a higher propensity of Chinese firms to enter sectors with already more incremental innovation would not be a threat to validity.

Note that for the sample period from 2001 to 2005 (and similarity measures in each of those years being compared to five years earlier), the assumption of the model that market size is fixed seems plausible from a European perspective. The Chinese domestic product market is likely much less important for European producers than access to the European market is for Chinese exporters during this time period.

Figure 2.6a and the kernel density estimation in Figure B.6 show that there is significant variation in (the five-year difference of) the similarity measure over the sample of patents. The sample used in the estimation is considerably smaller than the one used in Bloom et al. (2016) as only firms with patents in a given year and five years before can be used. Furthermore, not all firms have been successfully matched to patent data and, depending on the comparison

[28] In most of their estimations, Bloom et al. (2016) rely on an instrument using quotas in textiles and apparel that were lifted. As this restricts the sample to a sector that may not be representative and for which firms' ability to choose the type of innovation may be particularly limited by technological constraints, I use their second instrument. Appendix B.5.4 shows that the IV estimates are similar, but are likely to be biased towards OLS due to the weak instrument in this smaller sample. The F statistic is even well below the "less strict rule of thumb" of 5 (Cameron and Trivedi, 2005).

Figure 2.6: Histograms of Five-Year Change in Similarity Measure

(a) All patents of firms in sample

(b) Sample used in the regressions based on similarity to firms' own patents

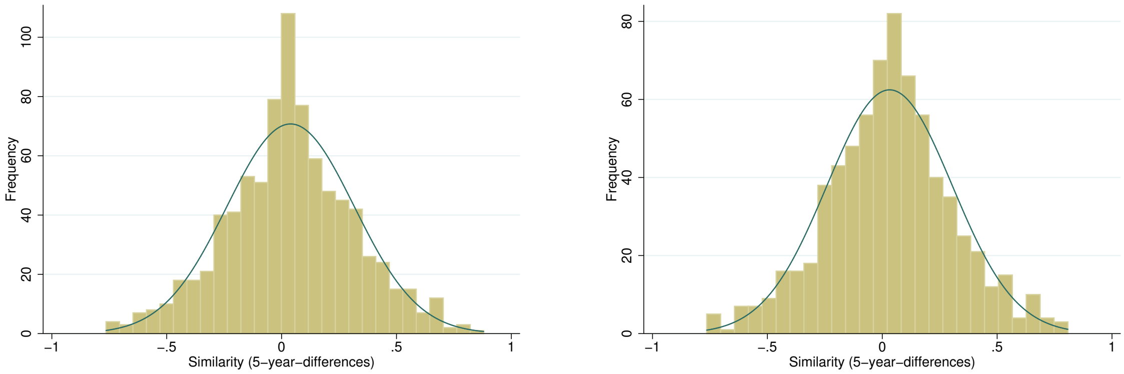

group, the firms' similarity measure may be missing for some years. For example, if the comparison group of patents consists of the firm's own patents in the preceding four-year period, firms which have not patented before will not be assigned a similarity score. The sample size therefore increases as the comparison group is defined more broadly at different levels of industry aggregation. The sample used in the main estimation with a comparison to a firm's own patents, which leads to the smallest sample, is shown in Figure 2.6b.

Table 2.1: Descriptive Statistics

	mean	sd	min	p10	p25	p50	p75	p90	max
Δ sim(comp.)	0.03	0.27	-0.76	-0.30	-0.13	0.04	0.19	0.38	0.81
Δ sim(SIC4)	0.02	0.17	-0.51	-0.20	-0.09	0.02	0.13	0.23	0.66
Δ sim(SIC3)	0.01	0.16	-0.60	-0.19	-0.09	0.02	0.11	0.20	0.52
Δ sim(SIC2)	0.01	0.15	-0.46	-0.17	-0.08	0.02	0.10	0.19	0.59
Δ sim(all)	0.01	0.15	-0.49	-0.17	-0.08	0.01	0.10	0.18	0.55
Δ ShareImChina	0.03	0.05	-0.06	0.00	0.00	0.01	0.03	0.06	0.54
$\text{ShareImChina}_{jt-6}$ $\cdot\Delta\text{ShareImChina}_t$	0.15	0.12	0.00	0.05	0.08	0.11	0.18	0.25	0.75
Observations	868								

Notes: The descriptives in this table include the observations of the estimation based on the comparison to all previous patents (see Table 2.2, column 5). This sample is the largest one of the five columns, since it is more likely that comparison group patents exist than for the more narrowly defined comparison groups. The sample includes five-year differences for the years from 2001 to 2005. The data are from PATSTAT, Bloom et al. (2016) and Peruzzi et al. (2014).

The results of the two-stage least squares (2SLS) estimations are reported in Table 2.2 and show that an increase in import competition leads to a shift towards more similar, i.e., more incremental, corporate innovation for a given firm. The coefficient in column 1 uses the outcome measure for the similarity of a firm's patents in one year to its own patents filed in the preceding four years. The strongest effect, both in terms of size and statistical significance, is found here. This means that firms in sectors with higher Chinese import penetration file patents that are closer to their past research agenda in terms of technological direction. It can be interpreted as a reduction in exploration or radical innovation as a result of more intense competition. In light of the model, this empirical result is in line with a decrease in the average size of the inventive step, such that new innovations are closer to previous ones in the technology space.

The next four columns in Table 2.2 use different similarity measures for the outcome (which is still the five-year difference in this measure). Columns 2 to 4 compare a firm's patents filed in a year to the entire industry's patents filed in the preceding four years, where the industry is defined on the four, three and two-digit SIC level, respectively. In the fifth column, the comparison group consists of all matched firms' patents in the previous four years. The coefficients of interest tend to become smaller as the comparison group is broadened to include more firms and varied sectors. While the sign remains positive, the effects in these four columns are not statistically significant. Hence, while there is clearly no evidence of a shift towards innovations becoming *more* different from any of the groups, the increase in similarity seems to be driven by a focus on a firm's own core research agenda.[29] Note that the sample contains only European firms, such that the new competitors from China are excluded. A conceivable reaction of European industries could be extended cooperation in R&D or even collusion. The absence of a significant effect on similarity to a firm's industry suggests that this is not the case. At the same time, additional specialization on niches could decrease similarity to the market, but does not seem to take place at a significant extent either.[30]

Again, if higher average similarity in an industry is positively correlated with competition, e.g. because technologies are more standardized and markets can be entered more easily by Chinese firms, this does not threaten the identification for two reasons. First, the model is

[29] An interesting area for future research is the relation between these findings and firm scope, as studied by Akcigit and Kerr (2010). It may also be that more intense competition provides incentives for narrower research and, as a result of the small steps remaining in the firm's original technical direction, its scope narrows.

[30] This conclusion does not change if a firm's own patents are excluded when the similarity to its industry is calculated. Note that my results are also consistent with Autor et al. (2016b), who do not find evidence for industry switching.

estimated in differences, such that only changes, but not differences in levels, are estimated. Second, a valid instrument ensures that what is estimated is indeed the effect of an increase in competition on the change of average similarity, rather than the reverse.

The larger coefficient in the first column is partly due to the larger average similarity of a firm's patents to its own previous research than to the other comparison groups. The descriptive statistics in Table 2.1 show that also the standard deviation is about twice as large for similarity to own patents as for the other similarities. However, dividing the coefficients in Table 2.2 by the standard deviations of the respective variables in the sample reveals that the effect remains larger for the comparison to own patents and decreases as the comparison group is widened.[31] One way to interpret the effect size is to calculate how a firm's position in the similarity distribution changes through import competition. The median of the similarity measure (and not of its change) for the sample in column 1 is 0.562. The mean of ΔShareImChina is 0.025. Multiplying this number by the coefficient 1.459 in column 1 and adding it to the median results in a similarity of 0.598. Thus the median firm is shifted by the mean effect from the 50th to the 56th percentile in the similarity distribution.

These results contribute the identification of a causal mechanism, which is playing a role in the trend away from risky basic research towards more applied and incremental R&D, to the observation by Arora et al. (2015). Venturing into the technological areas pertaining to other industries may yield more radical innovation through novel recombinations of ideas, but staying in one's field is likely to deliver more predictable results. When these latter small steps become relatively more valuable in the product market through increased risk of and opportunity for business stealing, firms adjust their research accordingly. Note also that the sample contains only firms which have patented in at least two years with a period of five years in between. Thus the subset of firms provides evidence of an actual change in these firms' behavior rather than a change in the sample of firms, e.g. through differences in survival, as considered in the next subsection.

[31] The coefficient for the change in import competition in Table 2.2, column 1, is 1.459. Divided by the standard deviation for the sample used in this regression, shown in the appendix in Table B.2 (which is the same as the one for the larger sample in Table 2.1), is 0.27 is 5.40. The same calculation for the coefficient in column 5 of Table 2.2 yields an effect size of (0.138/0.15 =) 0.92.

Table 2.2: Similarity

	(1)	(2)	(3)	(4)	(5)
	Δ sim(comp.)	Δ sim(SIC4)	Δ sim(SIC3)	Δ sim(SIC2)	Δ sim(all)
	b/se	b/se	b/se	b/se	b/se
Δ ShareImChina	1.459***	0.314	0.288	0.135	0.138
	(0.557)	(0.331)	(0.342)	(0.242)	(0.243)
Year FE	yes	yes	yes	yes	yes
Companies	298	421	425	428	428
SIC4	114	135	139	142	142
SIC3	67	74	75	78	78
SIC2	17	19	19	19	19
N	707	860	865	868	868

Notes: The dependent variable is the five-year difference in the similarity measure (based on SVD with 100 components), which uses the abstracts of a firm's patents in a year to calculate their similarity to the comparison group of patents in the preceding four-year period. In column 1, the comparison group consists of a firm's own past patents. In columns 2 to 4, the comparison group contains patents of all firms in the same industry, where the industry is defined increasingly broadly based on the Standard Industrial Classification (SIC) (using four, three and two-digit SIC categories, respectively). In column 5 the comparison group includes patents of firms from all other industries as well. The model includes year fixed effects (see Equation 2.10) and is estimated with two-stage least squares using Bloom et al.'s (2016) initial conditions variable to instrument for the increase in Chinese import competition. The sample period is from 2001 to 2005. Standard errors are clustered at the four-digit SIC level. The data are from PATSTAT, Bloom et al. (2016) and Peruzzi et al. (2014).
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

The smaller sample compared to Bloom et al. (2016) is a result of the data requirements for the similarity measure (and of the imperfect match between data sets). The similarity measure based on a firm's own patents can be calculated only if there are at least four patents in the data (two in the two years for which the five-year difference is calculated and two for their respective comparison periods). Thus the sample is not representative of the total population of firms. However, when studying innovation, it may be precisely these particularly innovative firms, which remain in the sample, whose behavior we are interested in. If one considers an economy's overall innovation, the most innovative firms will necessarily have contributed to it disproportionately. Thus, policy makers may also be particularly interested in the response of these firms' R&D to increasing competition. Nonetheless, in Appendix B.5.2 regressions are repeated with adjusted weighting, such that each firm receives the same weight. The main effect is slightly smaller, but still significant at the 10% level.

Tables B.3 and B.4 in the appendix show the distribution of the firms included in the regressions in column 1 of Table 2.2 over countries and industries, respectively. Neither a country nor a two-digit SIC category contain more than a quarter of all firms in the sample and the distribution over countries and industries appears representative, with larger shares for larger countries and with Machinery and Chemicals among the top categories.[32]

One potential problem in these main regressions may result from the inclusion of patents from various different patent offices. The process of European integration has affected intellectual property rights and, in particular, more and more patents are filed with the European Patent Office rather than national offices. Furthermore, the availability of abstracts in English can be affected by the route through which a patent is filed. Thus, a general shift towards English or towards a different way of filing patents may confound the results if it has been affected by Chinese import competition as well.[33] To address these concerns, in Appendix B.5.1 I show the same regressions using different samples of patents. Including only USPTO patents at the cost of a further reduced sample size leads to qualitatively similar results (Table B.5). This robustness check also ensures that patents filed with multiple patent offices with delay are

[32]The finer four-digit SIC categorization reveals that SIC 3499 "Fabricated Metal Products, Not Elsewhere Classified" and SIC 2834 "Pharmaceutical Preparations" have the highest shares with 5.0% and 4.4% of firms in the sample, respectively. All other four-digit SIC categories include less than 4% of firms each. Repeating the regressions with either German or British firms excluded does not change the results qualitatively.

[33]Note, however, that level differences should not affect the results in the specification in five-year differences. Only *changes* correlated with Chinese import competition would pose a problem to identification.

not driving changes in the similarity measure, since only the U.S. filing would be included.[34] Table B.6 goes in the opposite direction and increases the sample to include patents filed in China and others with non-English abstracts (while duplicates are still removed). The results are qualitatively and in terms of statistical significance the same as in the main regression in Table 2.2.

The main regressions do not control for the change in the number of patents. As the results of Bloom et al. (2016) demonstrate, this number is itself affected by the change in Chinese import competition. It is therefore likely a "bad control", in the sense that the total effect of competition on the type of innovation is correlated with a change in the number of patents, because the latter is influenced by competition as well (Angrist and Pischke, 2008). Thus a regression including this control variable may lead to an erroneous attribution of some of the change in similarity to the change in patenting, while this may merely be a correlation. It is also possible, however, that the growth of patenting indeed affects similarity (and this may, at least partly, mediate the effect of competition on similarity). Table B.1 in the appendix reports the results and shows that the coefficients for the competition variable hardly change at all. An increase in patenting is correlated with lower similarity, which would suggest that the additional patents are due to the firms' entering new research areas rather than merely an increase in the propensity to file patents to protect intellectual property.[35]

Another concern regarding the process of patenting could be a shift towards secrecy. If such a shift were induced by Chinese import competition, e.g. due to weaker intellectual property rights in China, it would have to affect large innovations more than incremental ones to explain the results without an actual change in R&D. Even if this were the case and companies

[34]Another way to ensure that later filings in other countries do not influence the measure is to limit the sample to include only priority filings. This smaller sample leads to a weak instrument problem, however, and does not permit a reliable estimation, although the estimated effect is of a similar magnitude. Interestingly, the first stage improves if the sample is restricted further to include only Amadeus-PATSTAT matches above a matching score of 0.995 (the F statistic is 6.36) and the magnitude of the effect increases to about 3.5, which is significant at the 10% level. The cutoff for the matching score presents a tradeoff between sample size and matching accuracy. Qualitatively, however, the results are robust to selecting different cutoffs and 0.995 in particular leads to similar results with the sample used in Table 2.2 as well.

[35]This result alleviates a concern that has been raised in conversations with scientists at an innovation consultancy: Firms may be more likely to split inventions into multiple patents under increased competition in order to create patent thickets. Note also that simply increasing the number of patents would not bias the results, since only the frequency of terms, but not their absolute number, affects the measure. A bias could result if the additional patents, into which a given invention is split, are filed in different years. If this were the case, however, a positive correlation between more patenting and higher similarity would be expected.

merely avoided filing patents for large innovations without R&D adjustments, however, a decrease in the number of patents would be expected, contradicting the evidence. Moreover, the results below suggest that an initial focus on more incremental research increases sales under import competition. Note that Arora et al. (2015) conclude that their observations are not the result of a change in publication practices either.[36]

Appendix B.5.5 demonstrates the robustness of the results to variations in the methods used for text analysis.

Change of Firm Behavior or Firm Selection?

While the theoretical model keeps the number of firms constant and assumes that only their endogenous choice of R&D investments in different types of innovation is affected by competition, an overall causal relationship between competition and smaller step innovation in markets might as well be caused by an effect on the survival probability for firms with different, but fixed, R&D allocations. The previous estimation has identified a change in innovation for surviving firms (and which have continued to file patents). Of course, the two effects are not mutually exclusive and it is therefore interesting to explore whether the selection effect seems to play a role as well, when European firms' competition with Chinese exporters intensifies. Following Bloom et al. (2016), I combine their data with my similarity measures to estimate the change in various outcomes for firms previously working on different types of innovation. The absence of the probability of survival itself as an outcome in Table 2.3 is due to the fact that almost all of the firms in the sample survived. The sample of firms used here is larger than in the previous estimations as only a measure for initial similarity is needed, but not for current similarity. Nonetheless, this requirement limits the sample to (at

[36]While Arora et al. (2015) only make the binary distinction between scientific publications and patents, their finding that Chinese import competition is associated with lower stock market valuation of a firm's journal publications is consistent with my finding on sales. Their results on patents, however, merely confirm the increase observed by Bloom et al. (2016), but cannot draw a finer distinction between equally protected intellectual property in patents. Another explanation for the trend away from scientific publications proposed by Arora et al. (2015) is a decrease in firm scope, making applied research relatively more valuable. While these changes may play a role in the general trend, my results do not suggest that they are the channel through which import competition affects R&D. A firm which has previously patented in three areas and, after an increase in competition, continues patenting in just one of them, would have a lower similarity measure than a firm continuing in the three areas. Also, I do not find a significant shift towards more similar research within firms' main industries, while if peripheral research areas were discontinued, this might be expected.

least initially) innovative firms with sufficient patents for the construction of the similarity measures.

Table 2.3 reports the results of estimating the following equation with 2SLS (where the additional instrumental variable for the interaction term is the interaction of the initial conditions instrument with initial similarity):

$$\begin{aligned}\Delta\text{outcome}_{ijkt} = {} & \beta_1 \Delta\text{ShareImChina}_{ijkt} + \beta_2(\text{similarity}_{ijkt-5} \cdot \Delta\text{ShareImChina}_{ijkt}) \\ & + \beta_3 \text{similarity}_{ijkt-5} + \Delta f_t + \varepsilon_{ijkt},\end{aligned} \tag{2.11}$$

where the dependent variable is the five-year log-difference of five economic outcomes for the firms in the sample: employment, patent filings, new products, sales and total factor productivity. The coefficient β_2 estimates the differential effect of an increase in the share of Chinese imports and firms with a higher initial similarity score (where own patents are the comparison group).

Despite the larger sample, most of the coefficients of interest (belonging to the interaction term) are insignificant. This suggests that the main effect of increased import competition is on existing firms' behavior rather than on the economic outcomes for firms with potentially hard to adjust R&D strategies.

In Table 2.3, the largest coefficient is on sales, which are positively affected under increased competition for firms whose initial research has been more incremental. This positive effect on sales is significant at the 10% level and consistent with the model: A firm whose research has been more incremental benefits from the increased private economic value after an increase in competition. The small steps allow the firm to more easily steal business from foreign competitors or, presumably more likely in the setting of Chinese import competition in Europe, enable it to better protect its market share.

Table 2.3 does not adjust for multiple hypothesis testing. Adjustments like Holm (1979) or Benjamini and Hochberg (1995) would lead to the conclusion that none of the null hypotheses, that the outcomes are not affected by the interaction term, could be rejected at conventional

Table 2.3: Initial Similarity and Various Economic Outcomes

	(1) Empl. b/se	(2) Pat. b/se	(3) New products b/se	(4) Sales b/se	(5) TFP b/se
Δ ShareImChina	0.873	0.899	-1.846	-3.690**	-1.415
	(1.315)	(1.728)	(1.422)	(1.542)	(1.584)
Similarity ·Δ ShareImChina	-1.447	1.843	0.523	6.288*	6.229
	(3.036)	(4.311)	(2.223)	(3.298)	(4.770)
Similarity	-0.030	0.083	-0.029	-0.186	-0.125
	(0.117)	(0.161)	(0.082)	(0.122)	(0.146)
Year FE	yes	yes	yes	yes	yes
Companies	659	591	807	489	243
SIC4	175	167	176	154	110
SIC3	92	91	89	84	62
SIC2	20	19	20	19	18
N	1524	1299	1945	1005	500

Notes: These regressions estimate the effect of initial R&D's similarity measure (based on a comparison to a firm's own patents) and its interaction with Chinese import competition on five different economic outcomes, again using the initial conditions instrumental variable in 2SLS estimation. The dependent variables in columns 1 to 5 are the five-year log-differences in employment, patent filings, new products, sales and total factor productivity, respectively. Standard errors are clustered at the four-digit SIC level. The data are from PATSTAT, Bloom et al. (2016) and Peruzzi et al. (2014).
*** p<0.01, ** p<0.05, * p<0.1

significance levels. Note, however, that these economic outcomes are not independent, such that the dependence structure of the p-values is also likely to be far from independence (Romano et al., 2010). For instance, the Bonferroni correction may thus be viewed as too conservative. In particular, this dependence of outcomes even holds mechanically for the similarity measures, which is why no adjustment is made in Table 2.2. Correcting for the five hypotheses associated with the different similarity measures, the effect of competition on similarity to a firm's own previous patents would still be significant at the 5% level with the aforementioned procedures or the Bonferroni method.

Innovation Type and Forward Citations

The preferred measure of the type of corporate innovation in this study is based on latent semantic analysis of patent abstracts. This measure is capable of capturing a patent's position in technology space more precisely than traditional technological classifications, as it adjusts to new words and concepts and can continuously measure a patent's closeness to these concepts. While such text-based measures are increasingly employed in the literature in general (e.g. for industry classification in Hoberg and Phillips, 2010, or to measure uncertainty in Baker et al., 2016) and in research using patent data in particular thanks to increasing computational power, their calculation and interpretation is often less straightforward than traditional measures based on citations. Much of the literature has captured the heterogeneity between patents, e.g. in terms of economic value, using the number of citations or more sophisticated measures based on the network of citations.

This section studies the effect of Chinese import competition on the simple and arguably most popular measure of a patent's characteristics, which is the number of times a patent is cited by later patents (i.e., forward citations). This measure has been found to be a good proxy of the value of patents (e.g. Harhoff et al., 1999).[37] However, this study is not about the differences in firms' ability to produce valuable innovations, but about the type of innovations they aim to produce, which in the model is understood as a choice variable that may be determined

[37] As mentioned above, while the similarity measure for a given patent application is fixed at the time of filing, the value of a patent and its forward citations may be affected by competition later on. Nonetheless, forward citations are likely to contain significant information about the quality of an invention and are frequently used in the literature as such.

by an R&D strategy that is adjusted to product market conditions. Therefore the outcome variable is the five-year difference in the position of a firm's average patent in the citation distribution. These variables can also be interpreted as changes in the share of a firm's patents at the respective positions in the citation distribution and they replace the outcome variable in Equation 2.10. The first dependent variable in Table 2.4 is an indicator which is one if a patent has never been cited and zero otherwise. The second dependent variable (in columns 3 and 4) is an indicator for patents which are approximately in the middle 50% of the citation distribution. The last two columns test whether a firm's average patent is more likely to be in the top quartile of citations when Chinese import competition increases. Table 2.4 shows that the increase in never cited patents as a result of increased Chinese import competition (again using the initial conditions IV) is highly significant. The negative effect on patents in the middle of the citation distribution is significant at the 5% level. Furthermore, the effect of competition on the chance of a firm's average patent being among the top 25% most cited is negative, although not statistically significant.

Overall, the signs of the coefficients are consistent with a shift towards lower value patents as a result of higher competition. Import competition leads firms to produce more patents that are never cited and fewer patents with many citations. This is in line with the idea that firms adjust their R&D strategy towards more incremental and lower risk research, leading more likely to patentable innovations, but producing fewer breakthroughs. Controlling for the initial number of patents does not change the results qualitatively in the even-numbered columns.

Table 2.4: Effect on Forward Citations of Patents

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ not cited	Δ not cited	Δ middle 50%	Δ middle 50%	Δ top 25% cit.	Δ top 25% cit.
	b/se	b/se	b/se	b/se	b/se	b/se
Δ ShareImChina	1.176***	1.182***	-0.664**	-0.661**	-0.512	-0.521
	(0.407)	(0.407)	(0.280)	(0.277)	(0.500)	(0.499)
Δ ln(patents+1)		-0.005		-0.002		0.007
		(0.014)		(0.010)		(0.012)
Year FE	yes	yes	yes	yes	yes	yes
Companies	571	571	571	571	571	571
SIC4	159	159	159	159	159	159
SIC3	83	83	83	83	83	83
SIC2	19	19	19	19	19	19
N	1171	1171	1171	1171	1171	1171

Notes: The dependent variables are the five-year differences in the position of a firm's average patent in the distribution of forward citations. All columns include year fixed effects and the even-numbered columns control for the five-year change in patenting. The first two columns estimate the change in never cited patents and columns 3 and 4 the change in patents in the middle 50% of the citation distribution. The last two columns show the change in patents in the top 25% of the distribution. The model is estimated with two-stage least squares using Bloom et al.'s (2016) initial conditions variable to instrument for the increase in Chinese import competition. Standard errors are clustered at the four-digit SIC level. The data are from PATSTAT, Bloom et al. (2016) and Peruzzi et al. (2014).
*** p<0.01, ** p<0.05, * p<0.1

2.4 Conclusion

This study explores the link between product market competition and the type of corporate innovation. The literature has suggested intensified competition due to globalization as a potential cause of the shift towards less risky innovation that has been observed over the last decades (Arora et al., 2015). Building on existing models of competition and innovation (Aghion et al., 2001, 2005), a fundamental difference in the incentives for drastic and non-drastic innovation is shown to imply a shift towards a reduced share of large steps as competition increases. A decrease in pre-innovation profits due to more intense competition increases incentives for both types of innovation. However, only for small innovations do post-innovation profits increase as well and thereby increase incentives to invest *more* than for large innovation. Post-innovation profits for large innovations are constant due to the fact that they lead to a monopoly position (following the definition for drastic innovation in Tirole, 1988), i.e., the technological difference between leader and follower is so large that competition policy, which makes it easier for firms to steal business from the other, has no effect, since the competitor no longer operates in the market. The endogenous choice of investments in these two different types of innovation technologies is shown to be compatible with an inverted-U relationship between competition and total innovation. The maximum absolute level of large (small) innovation will occur at a lower (higher) level of competition than the maximum of total innovation. The abstract parameter representing competition in the model may be interpreted as a reflection of various policies affecting competition, such as antitrust policy, standardization of technologies (e.g. chargers for devices such as mobile phones or electric cars from different manufacturers) or the extent to which barriers to international trade exist.

The empirical part of this chapter develops a novel measure of the type of innovation a firm produces. Using latent semantic analysis, a firm's patents are compared to its own previous patents, the industry's previous patents or all patents. Similarity to this previous research is interpreted as the extent to which the research is incremental, i.e., the higher the similarity measure for a firm's new patent filings, the lower the share of large step innovation. This measure is used to identify the causal effect of an increase in competition in an important, albeit particular, setting: the impact of Chinese import competition on European firms' R&D. Following Bloom et al. (2016), I use instrumental variable regressions to estimate the effect of a change in Chinese import competition on the change in the similarity measure of a firm. The shift towards more incremental innovation caused by an increase in competition is confirmed in the data. Furthermore, the interpretation of the model as an adjustment of firms' investment behavior, rather than of the types of firms that remain in the market, is supported by the empirical results. Effects on outcomes for firms with different initial R&D are mostly insignificant, but provide suggestive evidence for an increased private value of incremental innovations, as predicted by the model.

Further studies are needed to estimate the causal relationship between competition and the type of corporate R&D in additional settings. For example, an understanding of the role of mergers and acquisitions as well as standardization in this relationship would be relevant for policy makers.

Depending on the social value, and in particular the positive externalities, of small and large step innovation, the socially optimal policy may not maximize the number of innovations, but should attach different weight to each type. Future research may clarify the optimal policy by estimating differences in the appropriability of different types of innovation and in their social value. Combining novel measures of innovation type with stock market reactions to patent grants of assignees and of their competitors (Kogan et al., 2017) may be a promising way to estimate their private and social values.[38] Such an approach may also be able to identify differential changes in these values after changes in competition.

The evaluation of policies affecting competition with respect to innovation is usually limited to the amount of innovation. By revealing a distortion towards more incremental innovation with increasing competition, this study has uncovered a hidden cost of lowering product market prices in this way, that may be overlooked with traditional measures of innovation. Depending on the differences in externalities of different innovation types, this cost may be substantial. In addition, a comprehensive welfare analysis also needs to take into account that corporate R&D is not the only source of technological progress. Hence, the relative strengths of public and private organizations at producing different types of innovations and the extent of spillovers also have to be considered in the optimal policy.

[38] Note that stock market reactions can capture only the "social value" for firms, but exclude consumer surplus, such that additional proxy measures of broader externalities could be included, especially when estimating welfare effects. (The assignee of a patent is the company or person receiving ownership of the patent that is granted.)

Chapter 3

Antitrust, Patents, and Cumulative Innovation: Bell Labs and the 1956 Consent Decree*

3.1 Introduction

Innovation is a key driver of economic growth. One of the main instruments governments use to foster innovation is the patent system. A patent gives the right to exclude others from using the patented inventions in order to stimulate innovation. However, there is a growing concern that dominant companies might use patents strategically to deny potential entrants, often small technology-oriented start-ups, access to key technologies in an attempt to foreclose the market.[1] As start-ups are thought to generate more radical innovations than incumbents, market foreclosure may harm technological progress and economic growth (Baker, 2012).[2] To address this problem many critics call for antitrust policies as a remedy (Wu, 2012; Waller and

*This chapter is based on joint work with Martin Watzinger, Markus Nagler, and Monika Schnitzer.

[1] Derek Thompson, "America's Monopoly Problem", The Atlantic, October 2016; Robert B. Reich, "Big Tech Has Become Way Too Powerful," The New York Times, September 18, 2015, p. SR3; Michael Katz and Carl Shapiro "Breaking up Big Tech Would Harm Consumer," The New York Times, September 28, 2015, p. A24; Thomas Catan "When Patent, Antitrust Worlds Collide," Wall Street Journal, November 14, 2011.

[2] For example, Akcigit and Kerr (2010) show that start-ups do more explorative research and Foster et al. (2006) show that in the retail sector the fast pace of entry and exit is associated with productivity-enhancing creative destruction.

Sag, 2014). Yet, up to now there are no empirical studies showing that antitrust enforcement can effectively promote innovation.

In this study we investigate whether patents held by a dominant firm are harmful for follow-on innovation, and if so, whether antitrust enforcement in the form of compulsory licensing of patents provides an effective remedy. We advance on these questions by analyzing the effects of one of the most important antitrust rulings in U.S. history: The 1956 consent decree against the Bell System. This decree settled a seven-year old antitrust lawsuit that sought to break up the Bell System, the dominant provider of telecommunications services in the U.S., because it allegedly monopolized "the manufacture, distribution, and sale of telephones, telephone apparatus and equipment" (Antitrust Subcommittee, 1958, p.1668). Bell was charged with having foreclosed competitors from the market for telecommunications equipment because its operating companies had exclusive supply contracts with its manufacturing subsidiary Western Electric and because it used exclusionary practices such as the refusal to license its patents.

The consent decree contained two main remedies. The Bell System was obligated to license all its patents royalty-free and it was barred from entering any industry other than telecommunications. As a consequence, 7,820 patents or 1.3% of all unexpired U.S. patents in a wide range of fields became freely available in 1956. Most of these patents covered technologies from the Bell Laboratories (Bell Labs), the research subsidiary of the Bell System, arguably the most innovative industrial laboratory in the world at the time. The Bell Labs produced path-breaking innovations in telecommunications such as the cellular telephone technology or the first transatlantic telephone cable. But more than half of its patents were outside the field of telecommunications because of Bell's part in the war effort in World War II and its commitment to basic science. Researchers at Bell Labs are credited for the invention of the transistor, the solar cell, and the laser, among other things.

The Bell case is uniquely suited to investigate the effects of compulsory licensing as an antitrust measure for two reasons: First, it allows to study the effects of compulsory licensing without any confounding changes in the market structure. In compulsory licensing cases, antitrust authorities usually impose structural remedies such as divestitures, which makes it difficult to separate the innovation effects of changes in the market structure from the innovation effects of changes in the licensing regime. Yet, in the case of Bell no structural remedies were imposed, despite the original intent of the Department of Justice. This was due to the intense lobbying of the Department of Defense as Bell was considered vital for national defense purposes.

Second, Bell's broad patent portfolio enables us to measure the effect of compulsory licensing on follow-on innovation in different competitive settings. 42% of Bell's patents were related to the telecommunications industry. In this industry, Bell was a vertically integrated monopolist who allegedly foreclosed rivals. The remaining 58% of Bell's patent portfolio had its main application outside of telecommunications. In these industries, Bell was not an active market participant. By looking at the differential effects of compulsory licensing inside and outside of the telecommunications industry we can distinguish the effects of potential foreclosure of patents and of potential bargaining failures that are inherent in the patent system.

Our analysis shows that compulsory licensing increased follow-on innovation that builds on Bell patents. This effect is driven mainly by young and small companies. But the positive effects of compulsory licensing were restricted to industries other than the telecommunications equipment industry. This suggests that Bell continued to foreclose the telecommunications market even after the consent decree took effect. Thus, compulsory licensing without structural remedies appears to be an ineffective remedy for market foreclosure. The increase of follow-on innovation by small and young companies is in line with the hypothesis that patents held by a dominant firm are harmful for innovation because they can act as a barrier to entry for small and young companies who are less able to strike licensing deals than large firms (Lanjouw and Schankerman, 2004; Galasso, 2012). Compulsory licensing removed this barrier in markets outside the telecommunications industry, arguably unintentionally so. This fostered follow-on innovation by young and small companies and contributed to the long run technological progress in the U.S.

Looking at the results in more detail, we first consider the effect of compulsory licensing on innovations that build on Bell patents. We measure follow-on innovation by the number of patent citations Bell Labs patents received from other companies that patent in the U.S. We find that in the first five years follow-on innovation increased by 17% or a total of around 1,000 citations. Back-of-the-envelope calculations suggest that the additional patents other companies filed as a direct result of the consent decree had a value of up to $5.7 billion in today's dollars. More than two-thirds of the increase is driven by young and small companies and individual inventors unrelated to Bell. Start-ups and individual inventors increase follow-on innovation by 32% while for large and old companies the increase is only around 6%. Robustness checks show that the increase in follow-on innovation is not driven by simultaneous contemporary shocks to technologies in which Bell was active or by citation substitution.

The increase in follow-on innovation by other companies is accompanied by a decrease in follow-on innovation by Bell, but this negative effect is not large enough to dominate the positive effect on patenting by others. The limited negative response by Bell is most likely due to the fact that at the time of the consent decree, Bell was a regulated monopolist subject to rate of return regulation. Yet, the consent decree changed the direction of Bell's research. Bell shifted its research program to focus more on telecommunications research, the only business Bell was allowed to be active in.

In a second step we split the increase in follow-on innovation by industry. We do not find any increase in innovation in the telecommunications industry, the aim of the regulatory intervention. Compulsory licensing fostered innovation only outside of the telecommunications industry. This pattern is consistent with historical records that Bell continued to use exclusionary practices after the consent decree took effect and that these exclusionary practices impeded innovation (Wu, 2012). As no structural remedies were imposed Bell continued to control not only the production of telephone equipment but was - in the form of the Bell operating companies - also its own customer. This made competing with Bell in the telecommunications equipment market unattractive even after compulsory licensing facilitated access to Bell's technology. For example, the Bell operating companies refused to connect any telephone that was not produced by Western Electric, the manufacturing subsidary of the Bell System (Temin and Galambos, 1987, p.222). In other industries, compulsory licensing was effective to foster innovation by young and small companies since Bell as the supplier of technology did not control the product markets through vertical integration or via exclusive contracts.

Although the 1956 consent decree was not effective in ending market foreclosure, it permanently increased the scale of U.S. innovation. In the first five years alone, the number of patents increased by 25% in fields with compulsorily licensed patents compared to technologically similar fields without; and it continued to increase thereafter. This increase is again driven by small and new companies outside the telecommunications industry. We find only a small increase in patents related to the production of telecommunications equipment. This indicates that market foreclosure may slow down technological progress and suggests that antitrust enforcement can have an impact on the long-run rate of technological change. In an in-depth case study we also show that the antitrust lawsuit led to a quicker diffusion of the transistor technology, one of the few general purpose technologies of the post-World War II period.

We contribute to the literature by being the first to empirically investigate the effect of antitrust enforcement on innovation. Our results suggest that foreclosure impedes innovation and that compulsory licensing without structural remedies is not sufficient to overcome foreclosure. Access to technology through compulsory licensing alone does not stimulate market entry and innovation unless there is sufficient access to the product market as well. These insights are relevant not only for antitrust cases about abuse of a dominant market position, such as the Bell case, but also for merger and acquisition cases where compulsory licensing is often used as a remedy when mergers are approved. Our empirical findings support theoretical arguments in the antitrust literature suggesting that to increase innovation, antitrust measures should focus on exclusionary practices and the protection of start-ups (Segal and Whinston, 2007; Baker, 2012; Wu, 2012).

We also contribute to the literature on intellectual property by providing robust causal evidence for the negative effects of patents on follow-on innovation of small and young companies. Our estimate of an increase in follow-on innovation by 17% is significantly smaller than the increase reported by Galasso and Schankerman (2015b). They study the innovation effect of litigated and invalidated patents and find an increase of 50%.[3] While our study looks mainly at patents in the electronics and computer industry, Sampat and Williams (2015) consider gene patents and find no effect on follow-on research. The size of our measured effects is consistent with that reported by other studies such as Murray and Stern (2007) and Moser

[3]Litigated patents are selected by importance and by the virtue of having a challenger in court. Thus, the blocking effects of these particular patents might be larger than the average effect for the broad cross-section of patents.

and Voena (2012). They study various measures of follow-on innovation and report an overall impact of a patent removal of about 10-20% in biotech and chemistry. Our finding of entry of companies as the main mechanism driving the positive innovation effects of compulsory licensing is consistent with Galasso and Schankerman (2015b). They show that the increase in citations can be attributed to small companies citing invalidated patents of large companies.

Finally, this study contributes to our understanding of innovation and growth in the United States in the 20th century. By providing free state-of-the-art technology to all U.S. companies, compulsory licensing increased U.S. innovation because it opened up new markets for a large number of entrants. This interpretation is consistent with theoretical concepts and historical accounts. Acemoglu and Akcigit (2012) show theoretically that compulsory licensing can foster innovation because it enables more companies to compete for becoming the leader in an industry.[4] In line with this idea, Gordon Moore, the co-founder of Intel, stated that "One of the most important developments for the commercial semiconductor industry (...) was the antitrust suit filed against [the Bell System] in 1949 (...) which allowed the merchant semiconductor industry 'to really get started' in the United States (...) [T]here is a direct connection between the liberal licensing policies of Bell Labs and people such as Gordon Teal leaving Bell Labs to start Texas Instruments and William Shockley doing the same thing to start, with the support of Beckman Instruments, Shockley Semiconductor in Palo Alto. This (...) started the growth of Silicon Valley" (Wessner et al., 2001, p. 86) Similarly, Peter Grindley and David Teece opined that "[AT&T's licensing policy shaped by antitrust policy] remains one of the most unheralded contributions to economic development – possibly far exceeding the Marshall plan in terms of wealth generation it established abroad and in the United States" (Grindley and Teece, 1997).

The remainder of this chapter is organized as follows. Section 3.2 describes the antitrust lawsuit against Bell and the consent decree. In Section 3.3 we describe the data and the empirical strategy. In Section 3.4 we show that compulsory licensing increased follow-on innovation and conduct robustness checks. In Section 3.5 we examine the effectiveness of compulsory

[4] In the model of Acemoglu and Akcigit (2012), compulsory licensing also makes innovation less profitable because leaders are replaced more quickly. In the case of Bell, compulsory licensing was selectively applied to only one company which was not active in the newly created industries. This suggests that there was no disincentive effect and that our empirical set-up cleanly measures the effects of an increase in competition on innovation.

licensing as an antitrust measure against foreclosure in the market for telecommunications equipment. In Section 3.6, we present the long run effects of the consent decree on U.S. patenting. Section 3.7 presents a case study of the licensing of Bell's transistor technology, the defining general purpose technology of the 20th century. Section 3.8 concludes.

3.2 The Bell System and the Antitrust Lawsuit

In this section we describe the Bell System and the antitrust lawsuit against Bell. We then discuss the unique features of the case that make it ideally suited for our empirical analysis.

3.2.1 The Bell System was a Vertically Integrated Monopolist

In 1956, American Telephone & Telegraph (AT&T) was the dominant provider of telecommunications services in the U.S. Through its operating companies, it owned or controlled 98% of all the facilities providing long distance telephone services and 85% of all facilities providing short distance telephone services. These operating companies bought all of their equipment from Western Electric, the manufacturing subsidiary of AT&T. As a consequence, Western Electric had a market share in excess of 90% in the production of telecommunications equipment. Western Electric produced telecommunications equipment based on the research done by the Bell Laboratories, the research subsidiary of AT&T and Western Electric. All these companies together were known as the Bell System, stressing its complete vertical integration (Figure 3.1). In terms of assets, AT&T was by far the largest private corporation in the world in 1956, employing 598,000 people with an operating revenue of $2.9 billion or 1% of the U.S. GDP at the time (Antitrust Subcommittee, 1959, p.31).

The Bell System held patents on many key technologies in telecommunications, as well as a large number of patents in many other fields. Between 1940 and 1970, Bell filed on average ~543 patents or 1% of all U.S. patents each year (see Figure 3.2). More than 70% of the patents protected inventions of the Bell Laboratories (Bell Labs), arguably the most innovative industrial laboratory in the world at the time.

Figure 3.1: The Structure of the Bell System

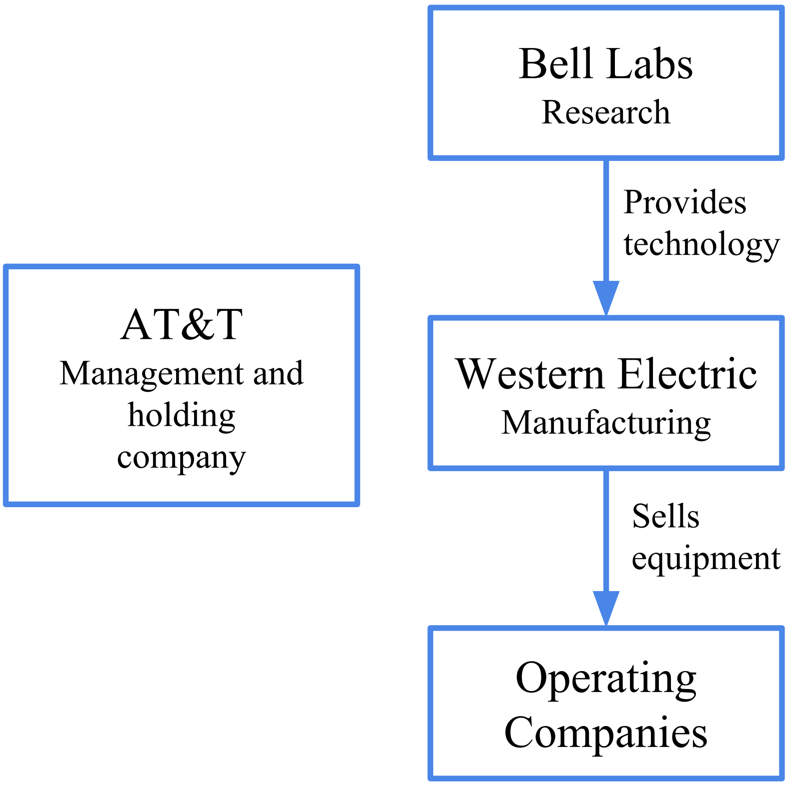

The Bell Labs were unique in their commitment to basic research. When the Bell Labs were founded in 1925, no one knew which part of science might yield insights for the problems of electric communication (Rosenberg, 1990; Nelson, 1962). As a result, the Bell System decided that - besides supporting the day-to-day need of the System - the Bell Labs would engage in basic science, assuming it would eventually yield products for some part of the large Bell System (Gertner, 2012; Nelson, 1959; Arora et al., 2015).[5]

The Bell Labs produced path-breaking basic and applied research. Scientists at Bell are credited for the development of radio astronomy (1932), the transistor (1947), cellular telephone technology (1947), information theory (1948), solar cells (1954), the laser (1957), and the Unix operating system (1969). The 1950 staff of Bell Labs alone consisted of four future Nobel Laureates in physics, one Turing Award winner, five future U.S. National Medals of Science

[5]According to the first head of basic and applied research at Bell Labs, Harold Arnold, his department would include "the field of physical and organical chemistry, of metallurgy, of magnetism, of electrical conduction, of radiation, of electronics, of acoustics, of phonetics, of optics, of mathematics, of mechanics, and even of physiology, of psychology and meteorology". This broad focus led to major advances in basic science, but also to a large number of unused patents. For example, an investigation of the FCC in 1934 reported that Bell owned or controlled 9,255 patents but actively used only 4,225 covered inventions (Antitrust Subcommittee, 1958, p.3842).

Figure 3.2: Size and Diversity of Bell's Patent Portfolio

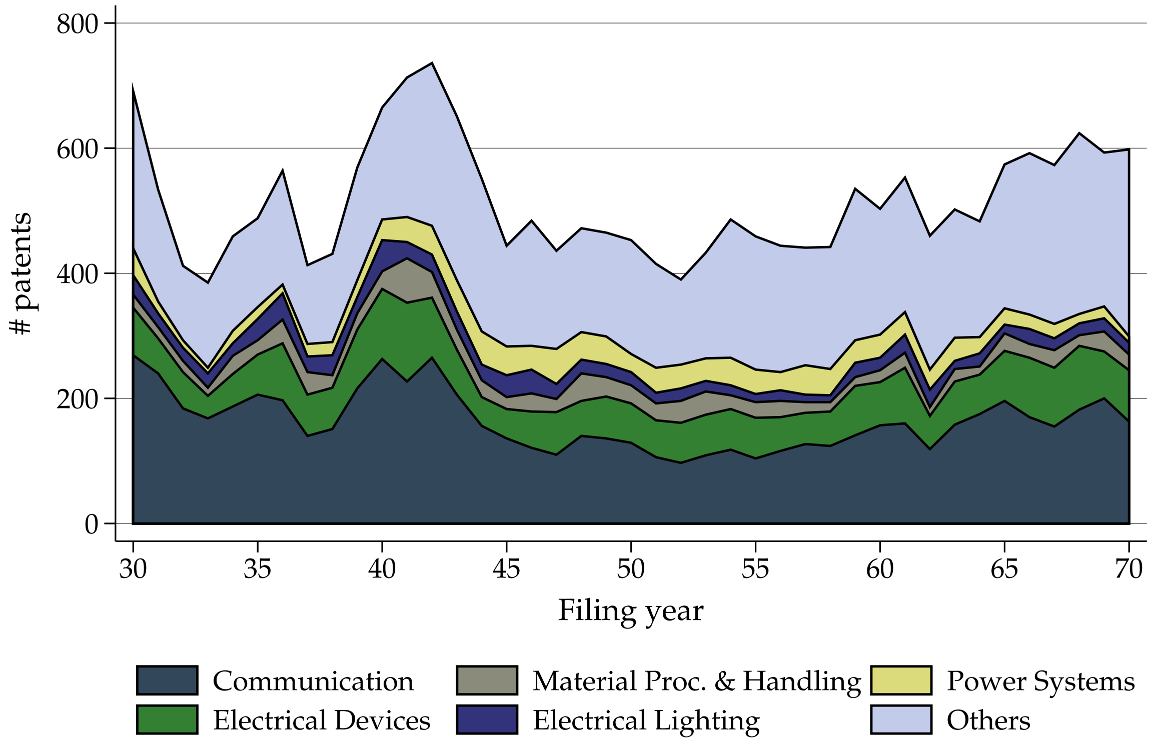

Notes: This figure shows the number of Bell patents in different technology subcategories over time. The subcategories aggregate the U.S. Patent Classification (USPC) following the scheme of Hall et al. (2001). We re-assign patents in the field of Optics to the Communication patents, as optics in the form of optical fiber played a large role in the development of communication technology starting in the 1960s. Of the compulsorily licensed patents published before January 1956 only 29 were in the field of Optics. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

recipients and 10 future IEEE Medals of Honor recipients. In 1950, Bell Labs employed 6,000 people, one third of whom were professional scientists and engineers (Nelson, 1962; Temin and Galambos, 1987). This was 1% of the entire science and engineering workforce in the U.S. at the time.[6]

3.2.2 The Antitrust Lawsuit

On January 14, 1949 the United States Government filed an antitrust lawsuit with the aim to split AT&T from Western Electric.[7] The complaint charged that Western Electric and

[6]According to the National Science Foundation, the number of workers in S&E occupations was 182,000 in the U.S. in 1950. Source: https://www.nsf.gov/statistics/seind12/c3/c3h.htm (last accessed 30 August 2016).

[7]This account of facts follows largely the final report to the Antitrust Subcommittee of the House on the Bell Consent Decree Program (Antitrust Subcommittee, 1959).

AT&T had been engaged in the monopolization of the manufacture, distribution and sale of telecommunications equipment in violation of the Sherman Antitrust Act of 1890 (Antitrust Subcommittee, 1959, p.46). According to the complaint, Bell was closing the market to all other buyers and sellers of telecommunications equipment by exclusionary practices including exclusive contracts and the refusal to license patents.[8]

To correct this, the government sought three main remedies. First, Western Electric was to be separated from AT&T, split into three competing companies, and to transfer all of its shares of the research subsidiary Bell Laboratories to AT&T. Second, AT&T was to buy telephone equipment only under competitive bidding and all exclusive contracts between AT&T and Western were to be prohibited. Third, the Bell System was to be forced to license all its patents for reasonable and non-discriminatory royalties (Antitrust Subcommittee, 1959, p.33).[9] Yet, none of this would happen.

The case ended with a consent decree on January 24, 1956, containing two remedies: First, the Bell System had to license all its patents issued prior to the decree royalty free to any applicant, with the exception of RCA, General Electric and Westinghouse who already had cross licensing agreements with Bell (the so called B-2 agreements). All subsequently published patents had to be licensed for reasonable royalties. As a consequence of the consent decree, 7,820 patents in 266 USPC technology classes and 35 technology subcategories (Figure C.1 in Appendix C.1) or 1.3% of all unexpired U.S. patents became freely available. Second, the Bell System was barred from engaging in any business other than telecommunications.

The decree was hailed by antitrust officials as a “major victory”, but already in 1957 the Antitrust Subcommittee of the Committee on the Judiciary House of Representatives started to investigate whether the decree of AT&T was in the public interest. The final report issued in 1959 pulled the decree to pieces: “the consent decree entered in the A.T. & T. case stands

[8]For example, Bell allegedly forced competitors “engaged in the rendition of telephone service to acquire AT&T patent license under threat of (...) patent infringement suits,” or refused “to issue patent licenses except on condition” to be able to control the telephone manufacturer or by “refusing to authorize the manufacture (...) of telephones (...) under patents controlled by (...) the Bell System” or by “refusing to make available to the telegraphy industry the basic patents on the vacuum tube” that are essential for telegraphy to compete with telephone or by refusing to purchase equipment “under patents which are not controlled by Western or AT&T, which are known to be superior” (Antitrust Subcommittee, 1958, p.3838).

[9]There were two minor remedies: First, AT&T was not to be allowed to direct the Bell operating companies which equipment to purchase and second, all contracts that eliminated or restrained competition were to be ceased.

revealed as devoid of merit and ineffective as an instrument to accomplish the purposes of the antitrust laws. The decree not only permits continued control by A.T. & T. of Western, it fails to limit Western's role as the exclusive supplier of equipment to the Bell System, thereby continuing monopoly in the telephone equipment manufacturing industry."

The hearings of the Senate subcommittee uncovered a timeline of cozy back and forth negotiations and intense lobbying by the Department of Defense (DoD). The DoD intervened on behalf of Bell because it relied on the research of the Bell Labs. In World War II, the Bell Labs had been instrumental in inventing the superior radar systems of the Allies. They also engaged in around a thousand different projects, from tank radio communications to enciphering machines for scrambling secret messages (Gertner, 2012, p.59 ff.).[10] In the following years, Bell Labs continued to work for the DoD, for example by operating the Sandia National Laboratories, one of the main development facilities for nuclear weapons.

After the complaint was filed in January 1949, Bell sought and obtained a freeze of the antitrust lawsuit in early 1952 with support of the DoD, on the grounds that Bell was necessary for the war effort in Korea. In January 1953, after Dwight D. Eisenhower took office, Bell began to lobby for the final dismissal of the case. The argument was that the Bell System was too important for national defense and thus should be kept intact. The government followed this argument and the Attorney General Herbert Brownell Jr. asked Bell to submit concessions "with no real injury" that would be acceptable in order to settle (Antitrust Subcommittee, 1959, p.55)

In May 1954, AT&T presented and in June 1954 submitted to the Department of Justice a checklist of concessions that would be an acceptable basis for a consent decree. The only suggested major remedy was the compulsory licensing of all Bell patents for reasonable royalties. To support its position, Charles Erwin Wilson, the Secretary of Defense, wrote Herbert Brownell Jr., the Attorney General, a memorandum to the effect that the severance of Western Electric from Bell would be "contrary to the vital interests of our nation" (Antitrust Subcommittee, 1959, p. 56). In December 1955, the Department of Justice communicated with AT&T that it was ready to consider a decree of the "general character suggested [by A. T.

[10] To highlight the engagement of Bell, we show in Figure C.2 in Appendix C.1.2 the patenting activity of Bell in radar and cryptography during World War II.

& T.] in its memorandum (...) dated June 4, 1954" (Antitrust Subcommittee, 1959, p.92). Bell agreed.

3.2.3 Advantages of the Bell Case for the Empirical Set-Up

The Bell case has two characteristics that make it ideally suited to measure the innovation effects of compulsory licensing as an antitrust remedy.

First, the consent decree did not impose any structural remedies for the telecommunications market. This allows us to isolate the innovation effect of compulsory licensing without any confounding changes in market structure. The reason why the Department of Justice did not impose any structural remedies is unclear. The final conclusion of the Antitrust Subcommittee blamed the lack of intent of the Attorney General to pursue Bell and the intense lobbying of the Department of Defense for the fact that no structural remedies were imposed (Antitrust Subcommittee, 1959, p.292). In contrast, the presiding judge Stanley N. Barnes stated that in his opinion it was enough to confine Bell to the regulated telecommunications market in order to prevent excessive prices and to end the exclusion of other suppliers (Antitrust Subcommittee, 1959, p.317).

Second, due to Bell Labs' commitment to basic science and its role in the war effort, Bell held a large number of patents unrelated to telecommunications, in industries in which it was not an active market participant. This gives us the opportunity to measure how the innovation effect of compulsory licensing depends on the market structure. In the telecommunications industry, Bell was vertically integrated. Hence Bell was not only a dominant player in the production of the technology used for telephone equipment, but it also controlled the production of telephone equipment (Western Electric), as well as the product market for telephone equipment through its operating companies. In all other industries, Bell was a supplier of technology, but was not active in production. Even more, the consent decree explicitly banned Bell from ever entering into these businesses which meant that it effectively preserved the market structure inside and outside of the telecommunications industry.

Figure 3.3: Compulsorily Licensed Patents by Industry

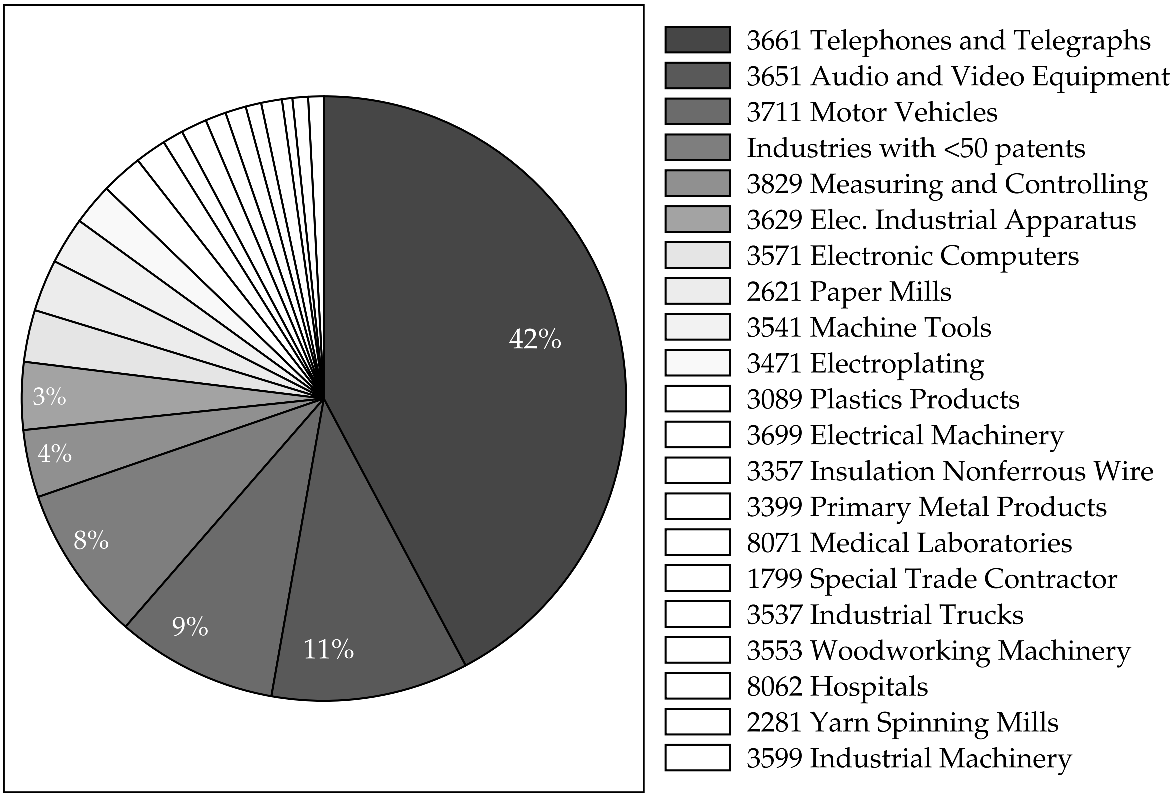

Notes: The pie chart shows the distribution of compulsorily licensed patents by most likely industry. We assign patents to the most likely four-digit SIC industry using the data of Kerr (2008). The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

To visualize the broad patent portfolio of Bell we use the data of Kerr (2008) to assign the most likely four-digit SIC industry group to each USPC class (Figure 3.3).[11] Around 42% of all Bell's patents have their most likely application in Bell's core business of producing telephones and telegraphs (SIC 3661). The remainder is spread across a large number of fields with an emphasis on electronics and industrial commercial machinery and computer equipment.[12]

3.3 Data and Empirical Strategy

For our estimation, we use comprehensive patent data for the U.S. from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office. In this data, we identify all

[11] We thank Bill Kerr for sharing his data.

[12] In Figure C.1 in Appendix C.1.1 we show the compulsorily licensed patents split by technology subcategories following Hall et al. (2001). Only 31% of all Bell patents are in the field of telecommunications and the remaining patents are spread over 34 other subcategories.

compulsorily licensed patents of the Bell System with a list of patent numbers published in the "Hearings before the Antitrust Subcommittee" of the U.S. Congress on the consent decree of Bell in May 1958 (Antitrust Subcommittee, 1958).[13]

In an ideal world, we would compare the number of realized follow-on innovations building on Bell patents with and without the consent decree. The problem is, however, that this is not possible: First, a census of follow-on innovations does not exist and second, we can observe only the state of the world in which the compulsory licensing of Bell patents happened but not the counterfactual situation without the consent decree.

To measure follow-on innovations we use patent citations. Bell patents could be freely licensed after the consent decree, but patents that built on licensed Bell patents still had to cite them. Thus, we can use patent citations as a measure for follow-on innovations even though patents had lost their power to exclude competitors (Williams, 2015). The advantage of this measure is that, in contrast to most alternative measures such as new products or R&D spending, citations are consistently available from 1947 onward.[14] Citations have the additional advantage that they have a high frequency which allows a precise measurement of effects. The caveat is that some citations might have been added by the patent examiner, which adds noise to the measure (Alcacer and Gittelman, 2006; Alcacer et al., 2009).

To construct a counterfactual for the compulsorily licensed Bell patents we use as control group all other patents that are published in the same year, that have the same total number of citations as the Bell patents in the five years prior to 1949, and that are in the same USPC technology class. By conditioning on the publication year and prior citations we control for the fact that, on average, young and high quality patents are cited more often. By conditioning on the same technology class we control for the number of companies that are active in the same field (i.e., for the number of potential follow-on inventors) and for technology-specific citation trends.

[13]The list is the complete list of all patents owned by the Bell System in January 1956. It also includes patents of Typesetter Corp. which were explicitly excluded from compulsory licensing in Section X of the consent decree. We assume that these patents are unaffected.

[14]In 1947 the USPTO started to publish citations of prior art on the front page of the patent (Alcacer et al., 2009).The first patent to include prior art was issued on February 4, 1947. Yet, inventions were evaluated against the prior art already since the passage of the Patent Act of 1836. Prior to 1947, however, the prior art was available only from the "file history" of the issued patent, which is not contained in PATSTAT.

We can interpret our results causally under the assumption that in the absence of the consent decree the Bell patents would have received the same number of citations as the control patents did (parallel trend assumption). More specifically, the identifying assumption is that conditioning on the control variables removes any systematic difference in follow-on citations between Bell and the control patents that is not due to compulsory licensing.

One potential concern about this identification strategy might be that the antitrust authorities chose to compulsorily license Bell patents for a reason related to the potential of follow-on research of these patents. According to the complaint and historical records, compulsory licensing was imposed because Bell used patents to block competitors in the field of telecommunications equipment. So if blocking patents are also patents that in the absence of compulsory licensing would have experienced particularly strong follow-on innovation then we might overestimate the effect of the consent decree.

Yet, this does not appear to be likely. In the absence of compulsory licensing, Bell's telecommunication patents would have continued to block competitors because the consent decree did not contain any other remedies aimed at restoring competition. Consequently, it seems fair to assume that blocking patents would have continued to receive the same number of citations as the control patents that have the same number of citations in the five years prior to 1949.

Furthermore, this concern obviously does not apply to the 58% of patents Bell held outside the field of telecommunications. These patents were included in the compulsory licensing regime of the consent decree not because they were blocking, but purely due to their association with the Bell System. Hence, there is no reason to expect any confounding effects.

To strengthen the point that the parallel trend assumption is plausible, we show in Section 3.4.1 that the number of citations of Bell and control patents was the same before the terms of the consent decree became known. In Section 3.4.3 we also show that companies that did not benefit from compulsory licensing did not start to cite Bell patents more after the consent decree. Thus, the control patents are a plausible counterfactual for patents both inside and outside of telecommunications.

Another concern might be that Bell's patenting strategy may have changed after the complaint became known. This is why we focus on patents *published* by 1949, the year the lawsuit against Bell started. The consent decree stated that only patents published before 1956 were to be compulsorily licensed. As a consequence of this cut-off date, more than 98% of the patents affected by the consent decree were filed before 1953, and more than 82% earlier than 1949. This implies that the characteristics of the majority of the affected patents were fixed before the Department of Justice filed its initial complaint. To be on the safe side, we use only patents granted before 1949, but the results do not change when we use all patents affected by the consent decree.

Out of the 7820 Bell patents affected by the consent decree, 4,731 patents were published before 1949. For 4,533 of these patents (i.e., for 95.8%) we find in total 70,180 control patents that fulfill the criteria specified above. In our empirical analysis, we use the weights of Iacus et al. (2009) to account for the potentially different number of control patents per Bell patent.[15]

Table 3.1 shows summary statistics. In column 1 we report the summary statistics for all patents published between 1939 to 1956. In column 2 we report the summary statistics of all Bell patents that were published between 1939 and 1956 and hence affected by the compulsory licensing rule. Patents published before 1939 had lost their patent protection by 1956 and were therefore not affected by the consent decree. In column 3 we report the summary statistics of the Bell patents published between 1939 and 1948. These are the patents that we use in our baseline regression.[16] They are affected by the consent decree but published before the lawsuit started and hence unaffected by a potential patenting policy change the lawsuit may have triggered.

The summary statistics of Bell patents differ from those of non-Bell patents. The average non-Bell patent in our data set receives 3.3 citations per patent and 6.1% of these citations are self-citations.[17] Bell System patents published in the same time period on average receive 5.2

[15] Iacus et al. (2009) proposes to use a weight of 1 for the treatment variable and a weight of $N_{Treatment,Strata}/N_{Control,Strata} \cdot N_{Control}/N_{Treatment}$ where $N_{Control}$ is the number of control patents in the sample, $N_{Control,Strata}$ is the number of control patents in a strata defined by the publication year, the USPC primary class and the number of citations up to 1949. $N_{Treatment}$ and $N_{Treatment,Strata}$ are defined analogously. Using these weights we arrive at an estimate for the average treatment effect on the treated.

[16] To make the statistics comparable for affected and not affected patents, we only consider technology classes in which Bell is active.

[17] In the main part of our study we only use citations by U.S. patents. In the appendix we run one regression with citations of patents filed in foreign jurisdictions.

Table 3.1: Summary Statistics

	(1) Non-Bell System	(2) Bell System Affected	(3) Bell System Baseline Sample
	mean	mean	mean
Filing Year	1944.5	1943.6	1940.6
Publication Year	1947.6	1946.5	1943.1
# Years in patent protection after 1956	8.6	7.5	4.1
Total cites	3.3	5.2	4.9
Citations by other companies	3.1	4.5	4.3
Self Citations	0.2	0.7	0.7
Citations by other companies prior to 1949	0.3	0.9	1.4
Observations	293578	7820	4731

Notes: The table reports the average filing and publication year, the *average* number of years until patent expiration and citation statistics for patents published between 1939 and 1956. Column 1 includes all patents of non-Bell System companies in technologies where a Bell System company published at least one patent. Column 2 includes all Bell patents published between 1939 and 1956. Column 3 includes all Bell patents published between 1939 and 1949, the baseline sample of most of our regressions. A citation is identified as a self-cite if the applicant of the cited and citing patent is the same or if both patents belong to the Bell System. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

citations and 13.4% of these citations are self-citations.[18] The numbers for the subsample of Bell patents published until 1949 are very similar. They receive on average 4.9 citations of which around 14.2% are self-citations.

3.4 Results: Compulsory Licensing Increased Follow-On Innovation

Prior to the consent decree, Bell licensed its patents to other companies at royalty rates of 1% - 6% of the net sales price. Lower rates applied if a cross-license was agreed upon (Antitrust Subcommittee, 1958, p. 2685). The consent decree lowered these rates to zero and made licensing available without having to enter into a bargaining process with Bell. In this section we estimate whether and if so by how much this compulsory licensing increased follow-on innovations.

[18] Except when explicitly mentioned in the text we correct for self-citations in all our regressions because we are mainly interested to which extent other companies built on Bell Labs patents.

Figure 3.4: Effect of Compulsory Licensing on Subsequent Citations

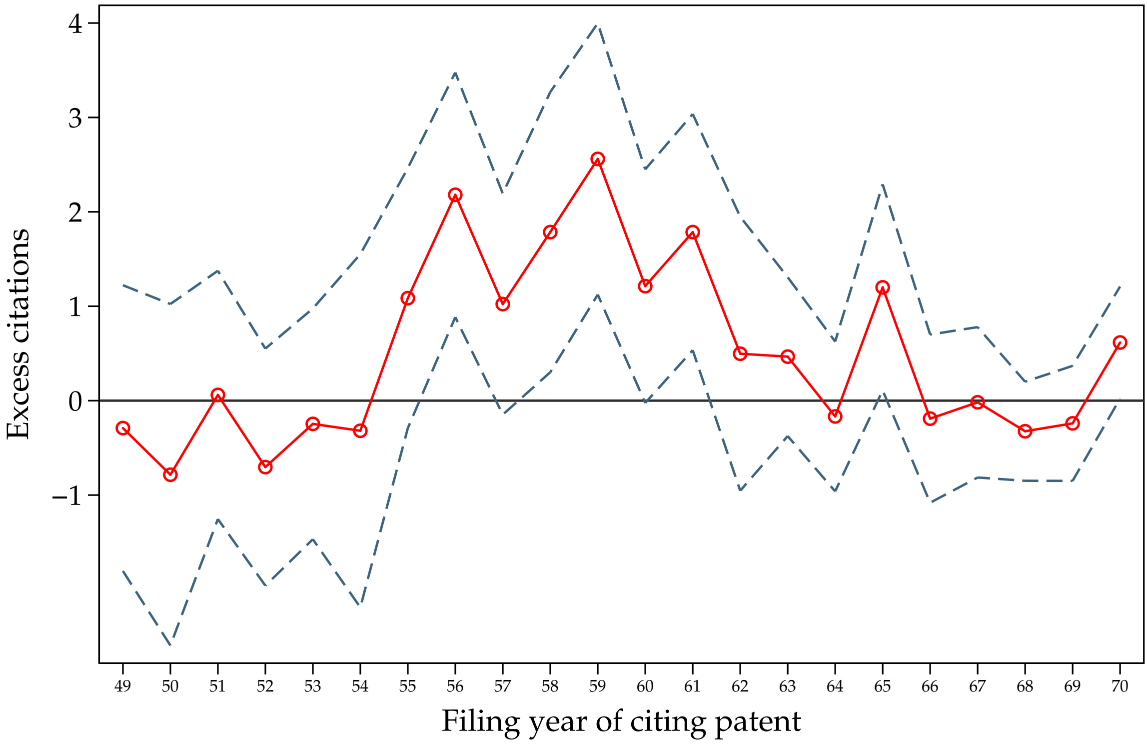

Notes: This graph shows the estimated number of yearly excess citations of patents affected by the consent decree ("Bell patents") relative to patents with the same publication year, in the same three-digit U. S. Patent Classification (USPC) primary class and with the same number of citations up to 1949. To arrive at these estimates we regress the number of citations in each year on an indicator variable that is equal to one if the patent under consideration is affected by the consent decree, and year fixed effects (Equation 3.1). We correct for self-citations. The dashed line represents the 90% confidence bands for the estimated coefficient. The sample under consideration contains 4,533 Bell patents and 70,180 control patents. We cannot match 198 Bell patents to control patents. To adjust for the different number of control patents per treatment patent in each stratum, we use the weights suggested by Iacus et al. (2009). The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

3.4.1 Timing: The Consent Decree Increased Citations of Other Companies Starting in 1955

In this section, we estimate the impact of the compulsory licensing on citations looking at the time period 1949-1970. We employ the following difference-in-differences specification:

$$\#Citations_{i,t} = \alpha + \beta_t \cdot Bell_i + YearFE_t + \varepsilon_{i,t} \tag{3.1}$$

where $\#Citations_{i,t}$ is the number of follow-on citations of other companies to patent i in year t. $Bell_i$ indicates whether the patent i is owned by the Bell System and is therefore treated. We also include fixed effects for each year ($YearFE_t$).

Figure 3.4 shows per year the estimated number of excess citations of Bell patents that were granted before 1949 relative to control patents, β_t in Equation 3.1. From 1949 to 1954, the average number of citations of treatment and control patents track each other very closely, speaking in favor of parallel trends in citations to Bell patents and to the control patents. In 1955, the average number of citations of other companies to Bell patents starts to increase and it converges again in 1960; 1960 is the average expiration date of the Bell patents in our sample (Table 3.1).[19] The yearly coefficients from 1955 to 1960 are mostly significantly different from zero at the 10 % level.[20]

The increase in citations depicted in Figure 3.4 does not start in 1956, the year of the consent decree, but in 1955. This is plausible because on May 28, 1954, Bell already suggested a consent decree including the compulsory licensing of Bell System patents as described in Section 3.2. Thus, both the Bell Laboratories and companies building on Bell's patents could have known that compulsory licensing was pending as early as May 1954 (Antitrust Subcommittee, 1959).[21]

This timeline is supported by the cumulative abnormal stock returns for AT&T stocks shown in Figure 3.5.[22] Up to the election of Dwight Eisenhower, cumulative abnormal returns were centered around zero. At the beginning of 1954, cumulative abnormal returns strongly increased to around 11%. The large uptick in March 1954 is exactly synchronized with the date of a memorandum summarizing a meeting of the Attorney General and Bell management about how to resolve the Bell case (Antitrust Subcommittee, 1958, p. 1956). Shortly thereafter, in May 1954, Bell proposed compulsory licensing as an acceptable remedy to settle the lawsuit. There is no more persistent positive or negative change in the cumulative abnormal return

[19] From 1861 to 1994, the term of the patent was 17 years from issuance.

[20] In Appendix C.2.1 we graphically compare the average yearly number of citations to Bell and to control patents and find the same results.

[21] The first media mentioning of the consent decree against Bell was on May 13, 1955 in the New York Times. Public officials confirmed that top level negotiations are ongoing "looking towards a settlement of the AT&T case".

[22] The historical stock market data is from CRSP.

Figure 3.5: Cumulative Abnormal Stock Returns of AT&T

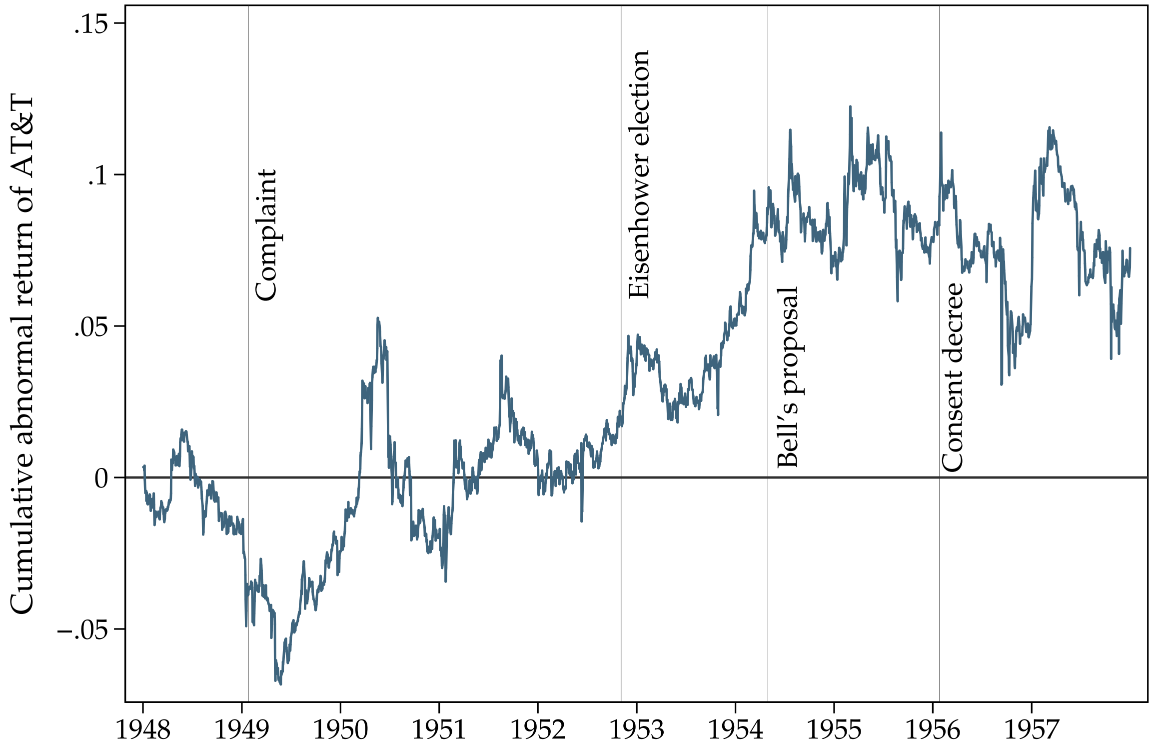

Notes: This figure shows the cumulative abnormal stock return of AT&T compared to other companies in the Dow Jones index, beginning in January 1948. The events marked in the graph are the beginning of the antitrust lawsuit on January 14, 1949, the presidential election on November 4, 1952, Bell's proposal of compulsory licensing on June 4, 1954, and the consent decree on January 25, 1956. The data are from the Center for Research in Security Prices (CRSP).

until 1959. In particular, the consent decree itself in 1956 did not seem to have had any more informational value.

We can also infer from Bell's behavior that as early as the first half of 1955, compulsory licensing was expected. According to the consent decree, all patents had to be licensed for free if they were published before January 24, 1956. If they were published after this cut-off date, they were licensed on a reasonable and non-discriminatory basis. So starting from the date when Bell became aware of the clause it had an incentive to delay the publication of its patents beyond the cut-off date.

According to the data, Bell indeed started to delay its patents at the patent office beginning in the first half of 1955. To pin down the date, we compare the propensity of a Bell patent to be published with the propensity that control patents are published for a given filing year. In

Figure 3.6, we show these hazard rates of publishing in a particular year for the filing years 1949 and 1953. For the filing year 1949, the publishing rates per year are very similar for Bell patents and patents from other companies. If at all, Bell patents were published a bit earlier. For the filing year 1953, this picture is reversed: Starting in the first half of 1955, Bell patents had a significantly lower probability of being published. This is consistent with Bell trying to delay the publications of its patents and having credible information about the general outline of the consent decree in the first half of 1955 at the latest.

Figure 3.6: Hazard Rates for Publication of Patents by Filing Year

(a) Filing year 1949

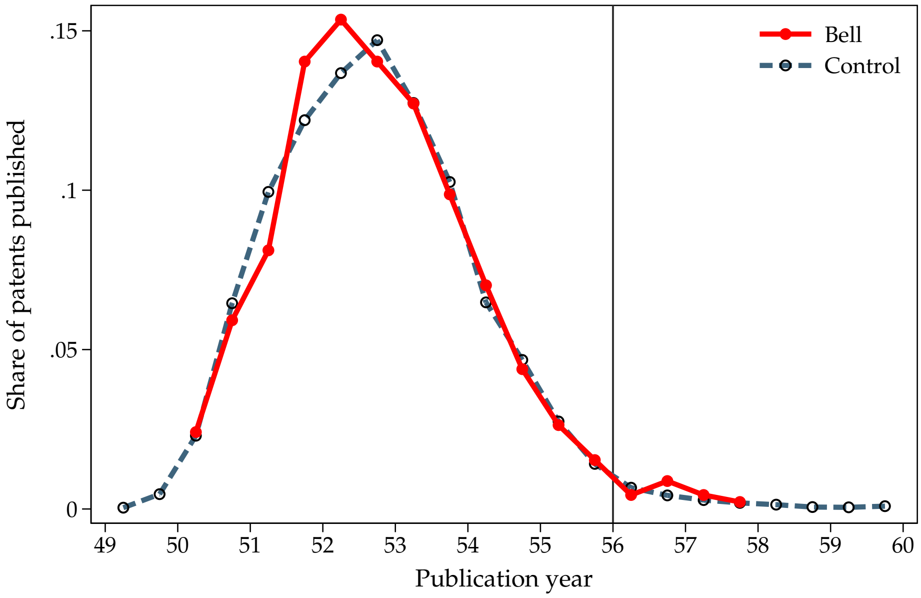

(b) Filing year 1953

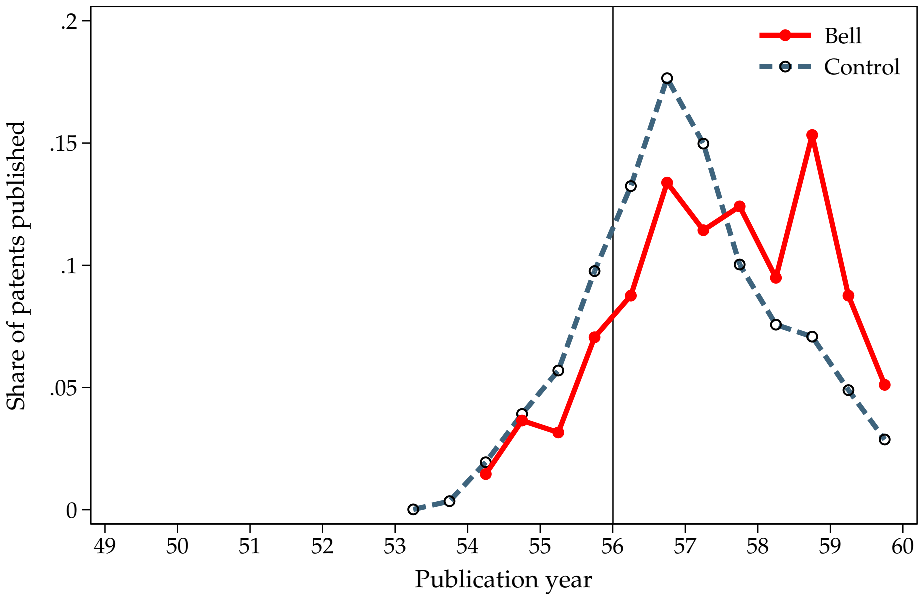

Notes: These figures show the hazard rates for publication of patents that were filed by Bell (solid line) and others (dotted line). Subfigure (a) shows hazard rates for patent applications filed in 1949, Subfigure (b) for applications filed in 1953. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

3.4.2 Magnitude: The Consent Decree Increased Citations to Bell Patents by 17%

We next present our baseline regression. To quantify the size of the effects of the consent decree, we estimate the average yearly effect of the consent decree on citations of other companies for the time period 1949-1960. We employ the following difference-in-differences model:

$$\#Citations_{i,t} = \beta_1 \cdot Bell_i + \beta_2 \cdot I[1955-1960] + \beta_3 \cdot Bell_i \cdot I[1955-1960] + \varepsilon_{i,t} \qquad (3.2)$$

where $I[1955-1960]$ is an indicator variable for the treatment period. We define the treatment period as from 1955 to 1960 based on the yearly coefficients in Figure 3.4.

The results are reported in Table 3.2 column 1.[23] In the treatment period, the consent decree resulted in 0.020 additional citations. This implies that, on average, the consent decree increased citations to Bell patents by other companies by 17% from 1955 to 1960.[24] Considering only the 4,731 patents published before 1949, this implies a total increase of 568 citations. If we assume homogeneous effects for all 7,820 patents published up to 1956, the total number of excess citations is 938. The effect is also positive and statistically significant if we include all patents up to 1956, the year of the consent decree (column 2).

Back-of-the-envelope calculations suggest that the additional patents for other companies directly induced by the consent decree had a total value of up to \$5.7 billion. To calculate this number we use estimates for the average dollar value derived from Kogan et al. (2017) to weigh each citing patent.[25] According to these estimates, each compulsorily licensed patent created an additional value of \$121,000 annually in the treatment period (column 3). Assuming homogeneous effects for all 7,820 patents in the treatment group, the consent

[23] Note that patents receive fewer citations post treatment because older patents in general receive fewer citations than younger patents. See Figure C.3 in Appendix C.2.1

[24] To determine the percentage increase, we first calculate the number of citations Bell patents would have received in the absence of the treatment (counterfactual), using the coefficients in Table 3.2 column 1. The counterfactual is 0.115 (= 0.183 − 0.004 − 0.064). We then divide the treatment effect, 0.02, by the counterfactual (0.02/0.115 = 0.174).

[25] Kogan et al. (2017) measure the value of a patent using abnormal stock returns around the publishing date of the patent. We use this data to calculate the average dollar value for a patent in each technology class and publication year.

decree led to around \$5.7 billion in economic value over six years, between 1954 and 1960. These calculations represent an upper bound because they assume that without the additional citations induced by the consent decree the patent would not have been invented (i.e., that the compulsorily licensed patent was strictly necessary for the citing invention).

The effect is measurable across the quality distribution of patents. We split all patents by the number of citations a patent received in the first five years after publication and present results in columns 4 and 5 of Table 3.2. We define a high-quality patent as a patent with at least one citation before 1949 and a low-quality patent as a patent with no citations. The effect is stronger for high quality patents, but the effect is also statistically significantly different from zero for low quality patents. The effect is also not exclusively driven by the computer industry, which was just about to start in 1956. In column 6, we report results when dropping all 491 Bell patents classified in the technology subcategories "Computer Hardware and Software", "Computer Peripherals" and "Information Storage" or "Others" (Hall et al., 2001) and find a similar effect. The effect is also not driven by the concurrent consent decrees of IBM in 1956 or RCA in 1958. IBM and RCA were defendants in an antitrust case with compulsory licensing as the outcome. We drop all citations from patents that also cite either the patents of RCA or the patents of IBM and report the results in column 7.

Table 3.2: The Effect of Compulsory Licensing on Subsequent Citations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Citations by								
	other companies							Bell	all
	Base-line	Up to 1956	Dollar weighted	Low quality	High quality	w/o Computer	w/o IBM & RCA	Self-Cites	Total-Cites
Bell	-0.4	-0.6***	0.1	-0.5	-0.3	-0.5	-0.6	1.4***	0.5
	(0.5)	(0.2)	(2.9)	(0.4)	(0.9)	(0.5)	(0.5)	(0.3)	(0.7)
I(55-60)	-6.4***	-3.3***	3.7	-0.5**	-12.4***	-6.9***	-6.2***	-1.0***	-6.8***
	(0.6)	(0.8)	(2.3)	(0.2)	(0.9)	(0.7)	(0.6)	(0.1)	(0.7)
Bell x I(55-60)	2.0***	1.9***	12.1***	1.0**	3.1***	2.2***	2.0***	-0.6**	1.6*
	(0.6)	(0.6)	(4.0)	(0.4)	(1.0)	(0.6)	(0.6)	(0.3)	(0.8)
Constant	18.3***	19.9***	83.2***	8.4***	28.2***	18.7***	17.6***	1.5***	19.0***
	(1.2)	(1.6)	(3.7)	(0.4)	(1.4)	(1.4)	(1.1)	(0.1)	(1.2)
# treated	4533	7111	4533	2279	2254	4042	4533	4444	4731
Clusters	225	253	225	194	179	160	225	223	223
Obs.	896556	1121648	896556	580356	316200	700500	896556	854592	828876

Notes: This table shows the results from a difference-in-differences estimation with 1949-1954 as pre-treatment period and 1955-1960 as treatment period. The estimation equation is:

$$\#Citations_{i,t} = \beta_1 \cdot Bell_i + \beta_2 \cdot I[1955-1960] + \beta_3 \cdot Bell_i \cdot I[1955-1960] + \varepsilon_{i,t}$$

where $I[1955-1960]$ is an indicator variable for the treatment period from 1955 to 1960. *Bell* is an indicator variable equal to one if a patent is published by a Bell System company before 1949. As control patents we use all patents that were published in the U.S., matched by publication year, primary USPC technology class, and the number of citations up to 1949. To adjust for the different number of control patents per treatment patent in each stratum, we use the weights suggested by Iacus et al. (2009). As dependent variable, we use all citations by companies other than the filing company in columns 1 through 7. In the second column, we extend our sample of affected patents to 1956, and in the third column we use the sample up to 1949 and weight each citation by the average dollar value of a patent in the same publication year and technology class derived from the values provided by Kogan et al. (2017). In columns 4 and 5, we split the sample by their citations prior to 1955 to measure quality of patents. A patent is classified as "high quality" if it has at least one citation prior to 1955 and it is classified as "low quality" otherwise. In column 6 we exclude patents that are classified in technology subcategories related to the computer, and in column 7, we exclude all citations by patents of IBM and RCA and all patents that cite IBM and RCA patents. IBM had a consent decree with a compulsory licensing of patents in 1956 as well and RCA had a consent decree in 1958. In column 8, we match patents on publication year, technology class and self-citations prior to 1949 and use self-citations as the dependent variable. In column 9 we match on total citations prior to 1949 and use total citations as outcomes. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office. All coefficients are multiplied by 100 for better readability. Standard errors are clustered on the primary three-digit USPC technology class level. *, **, and *** denote statistical significance at the 10%, 5% and 1% levels, respectively.

3.4.3 Robustness Check: No Increase in Citations by Untreated Companies

One concern for the estimation is that the effect of compulsory licensing on subsequent citations might be driven by a shock that increased follow-on innovation to Bell patents and was correlated with the consent decree. For example, the antitrust prosecutors might have chosen to press for compulsory licensing because they expected that there would be many follow-on innovations based on the high quality of the Bell's patents.

To see whether this might have been the case we analyze the citation patterns of unaffected companies to Bell patents and to the control patents. The 1956 consent decree singled out three companies that were explicitly excluded from the free compulsory licensing of Bell patents: the General Electric Company, Radio Corporation of America, and Westinghouse Electric Corporation. The reason was that these companies already had a general cross-licensing agreement, the "B-2 agreements" dated July 1, 1932. A fourth company, the International Telephone and Telegraph Company (ITT), was also not affected by the decree as it had a patent pool with Bell.

We repeat our baseline analysis but use only the citations of the B-2 companies (including ITT) as the dependent variable and report the results in Figure 3.7 and column 2 of Table C.1 in Appendix C.2.2. We do not find any effect. This suggests that the consent decree did not change the citation behavior of excluded companies and the measured effects are not due to a common technology shock. As these companies in total make up 12% of all citations to Bell patents, this null effect is not due to a lack of measurability.[26]

A second concern might be that due to the free availability of Bell technology, companies substituted away from other, potentially more expensive technologies. In Appendix C.2.3 we show the results of additional auxiliary analyses suggesting that the effects are not driven by citation substitution.

[26] We repeat our analysis also for foreign companies, which could also use Bell patents for free but which did not receive technical assistance, and report the results in Table C.1, column 3 in Appendix C.2.2. Similarly, we repeat our analysis for companies that already had a licensing agreement in place and compare them with companies without a licensing agreement (Table C.1, columns 4 and 5, Appendix C.2.2). As expected, we find that the effects are smaller for firms that were less affected by the consent decree.

Figure 3.7: Effect of Compulsory Licensing on Subsequent Citations Among Companies Exempt from the Consent Decree

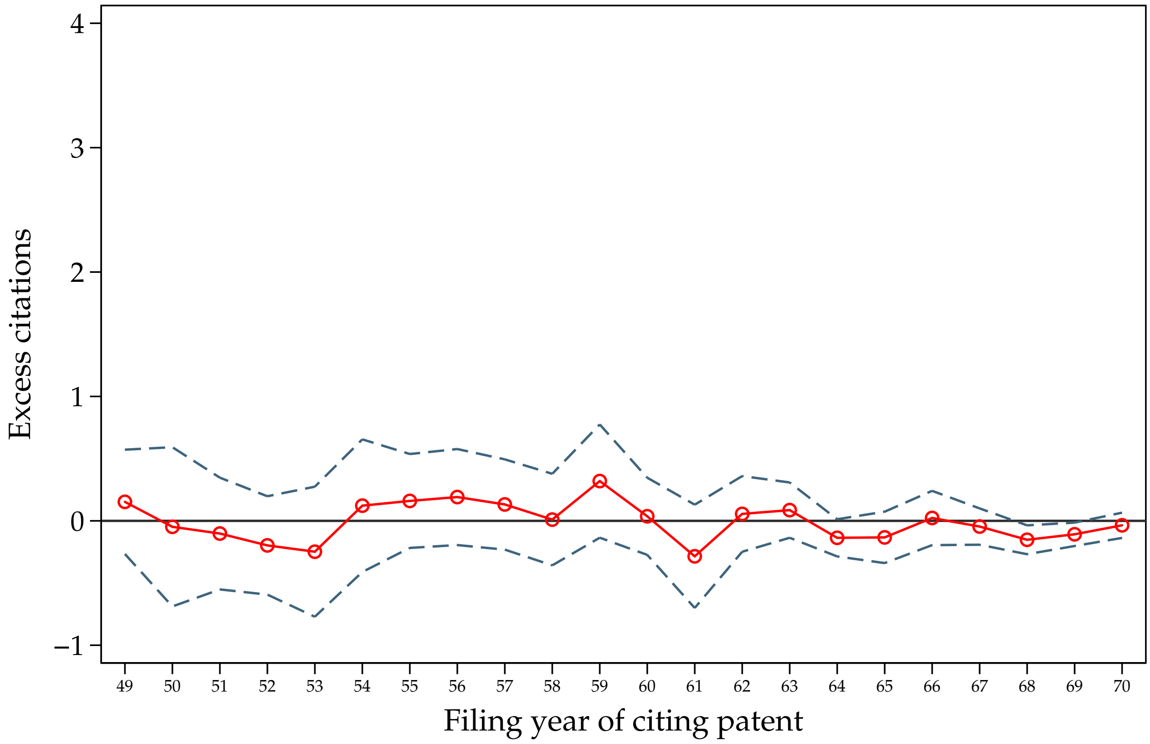

Notes: This graph shows the estimated number of yearly excess citations by General Electric Company, Radio Corporation of America and Westinghouse Electric Corporation, the three companies exempt from the consent decree, and by International Telephone and Telegraph Company, which already had a patent pool in place, of patents affected ("Bell patents") relative to patents with the same publication year, in the same three-digit USPC primary class and with the same number of citations up to 1949. To arrive at these estimates, we regress the number of citations by the unaffected companies in each year on an indicator variable equal to one if the patent under consideration is affected by the consent decree and year fixed effects. The dashed line represents the 90% confidence bands for the estimated coefficient. To adjust for the different number of control patents per treatment patent in each stratum, we use the weights suggested by Iacus et al. (2009). The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

Finally, in Appendix C.2.4 we vary the construction of control groups and show that our results are not driven by the particular choice of matching variables.

3.4.4 Robustness Check: The Decrease in Bell's Own Patenting is Lower than the Increase in Patenting by Other Companies

We next examine how Bell reacted to the consent decree. Bell might have reduced its innovation activities by more than other companies increased their innovation activities, such that the net effect of the consent decree would be negative. To see, whether this is the case we measure whether Bell continued to produce follow-on innovations building on its own patents.[27] Results are reported in column 8 of Table 3.2. The number of self-citations shows a decrease of 0.006 self-citations in the years between 1955 and 1960. This decrease is statistically significant, but is not large enough to dominate the increase in citations by other companies. In column 9 we present the effect on total citations, i.e., citations by other companies and self-citations by Bell. We find that total citations increased by 0.016. This speaks in favor of a net increase in innovation due to the consent decree.

Bell's innovation output in terms of number of patents continued to grow in line with expectations in the years following the consent decree. To show this, we construct a synthetic Bell and compare it with the actual patent output of the Bell System. To construct a synthetic Bell, we first calculate the share of Bell's patents of all patents in each technology subcategory for the years 1946, 1947, and 1948. Then we assume that Bell's growth would have been in line with the growth of other companies that existed before 1949 in these technology subcategories so that Bell would have held its share in each subcategory constant for the following years. Results are presented in Figure 3.8a. It shows that Bell's patenting is on average smaller than the patenting of the synthetic control, but not by much.[28]

Bell's continued investment in research was in line with the incentives the consent decree and the regulators provided. The consent decree did not significantly alter the profitability of new patents. The consent decree mandated that Bell could demand "reasonable" licensing fees for all patents published after January 1956. The reasonable royalty rates Bell charged

[27] Self-citations are a measure for how much a company develops its own patents further (Akcigit and Kerr, 2010; Galasso and Schankerman, 2015a).

[28] In Figure C.5 in the Appendix C.2.5 we compare the patenting output of Bell with other control companies and find that Bell's patent growth is in line - but at the lower end - of similar companies. The only exception is the growth of General Electric which is much larger, highlighting the problem of constructing a counterfactual for a single company.

Figure 3.8: Innovation and R&D in the Bell System After the Consent Decree

(a) Patenting Over Time: Bell System and Synthetic Bell

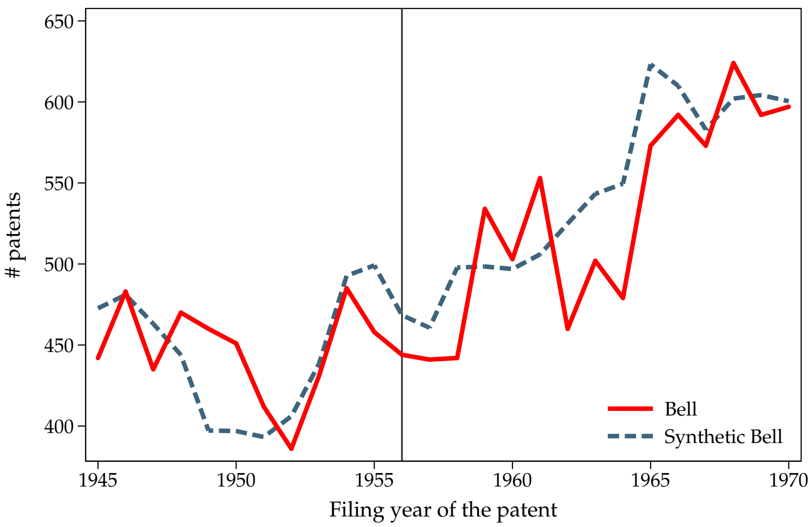

(b) R&D Spending by AT&T

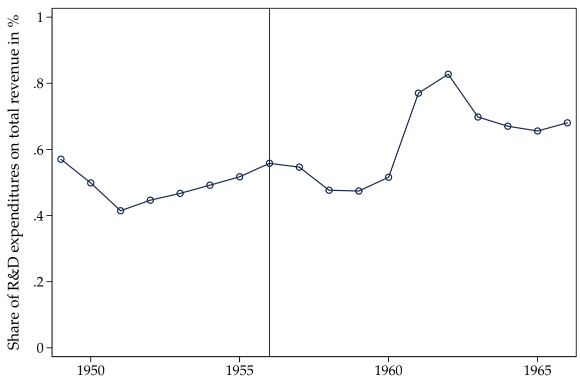

(c) Share of Communication Patents

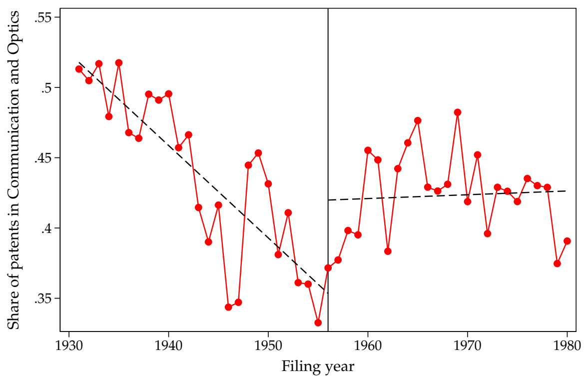

Notes: Subfigure (a) shows the total number of patents filed by the Bell System compared to a synthetic Bell. To construct the synthetic Bell, we calculate the share Bell's patents had in each two-digit technology subcategory relative to all patents of companies that had at least one patent before 1949. We then assume that in the absence of the consent decree, Bell's patenting would have grown in each subcategory at the same pace as the patenting of all other companies. As a consequence, Bell's share in each technology subcategory is held constant. In a last step, we add the number of patents up to a yearly sum. Subfigure (b) shows the ratio of R&D expenditures relative to total R&D of American Telephone & Telegraph. The data are from the annual reports of AT&T. Subfigure (c) shows the share of patents related to communication relative to all patents filed by Bell. We define a patent as related to communication if the most likely application is in the production of telecommunications equipment (SIC 3661). In Appendix C.2.6 we show the change in direction using NBER subcategories. The patent data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

were not much different compared to the pre-decree royalties (Antitrust Subcommittee, 1959, p.111). The only difference was that Bell had to give a license to any applicant.

Bell also had little incentive to reduce investment in R&D because the Bell System was subject to a rate of return regulation following the Communications Act of 1934. According to annual reports, AT&T had a stable ratio of R&D to operating revenue of 0.5% from 1949 to 1960 (Figure 3.8b).[29] For the entire Bell System, the share of R&D to total turnover stayed almost constant at 2%-3% from 1966 to 1982 (Noll, 1987). However, the absolute level of R&D effort increased as the Bell System grew. Operating revenues increased from $3.2 billion in 1950 to $5.3 billion in 1955, to $7.3 billion in 1960 and to $11 billion in 1965, while the staff at Bell Labs grew from 6,000 in 1950, to 10,000 in 1955, to 12,000 in 1960 and 15,000 in 1965 (Temin and Galambos, 1987).

But even if the consent decree offered no incentive for Bell to downsize, it offered incentives for Bell to redirect its research budget towards applications in the telecommunications field. Prior to the consent decree, Bell could expand to other businesses. Afterward, Bell's future was bound to common carrier telecommunications. The company correspondingly refocused its research program on its core business and increased its share of patents in fields related to the production of telecommunications equipment (Figure 3.8c).

These results are consistent with the study of Galasso and Schankerman (2015a) on patent invalidations. They show that large companies on average do not reduce follow-on innovations significantly if they lose a patent due to litigation. The only exception is if the large company loses a patent outside of its core-fields. Then it reduces innovation in the field of the patent under consideration and reacts by redirecting future innovation to a different but related field.

3.4.5 Mechanism: Increase in Citations is Driven by Start-ups.

We next examine which type of company increases innovation after the compulsory licensing and report the results in Table 3.3. We split citations by the type of the citing assignee. An assignee is either a company or an individual inventor; an assignee is defined as young and

[29] We do not know whether the consolidated balance sheet also includes the Bell Laboratories and Western Electric. It seems that at least some parts of the Bell System are not consolidated in the annual reports of AT&T.

small if its first patent was filed less than 10 years before it cited the Bell patent and if it had less than 10 patents before 1949.[30] We first use the number of citations from young and small assignees as the dependent variable and report the results in column 2. We then use the citations of all other assignees that are not young and small and report the results in column 3. In column 4 we look explicitly at small and young assignees that are companies ("start-ups"), leaving out individual inventors.[31]

We find that the increase in follow-on innovation is predominantly driven by young and small companies entering the market and by individual inventors. Young and small assignees increase their citations after 1955 by an average of 0.014 citations (32%) while all others increase their citations by 0.006 (6%) on average. Around 70% of the overall increase comes from young and small assignees, but they are responsible for only one-third of all citations to Bell patents (columns 2 and 3 in Table 3.3).[32] Among the small and young assignees, start-ups experience a particularly strong increase: they account for 50% of the total increase in citations although they are responsible for only 18% of all citations (column 4).

These results suggest that patents act as a barrier to entry for start-ups and prevent their follow-on innovation. They provide support for the hypothesis that the consent decree reduced potential bargaining failures. Several prior studies suggest that small firms might not have large enough patent portfolios to resolve disputes or to strike cross-licensing agreements (Lanjouw and Schankerman, 2004; Galasso, 2012; Galasso and Schankerman, 2015b). As cross-licensing was a priority in the licensing strategy of Bell prior to the consent decree, a small patent portfolio might have been a significant handicap for small inventors seeking a license from Bell (Antitrust Subcommittee, 1958, p. 2685).

One potential concern might be that the observed increase of citations by young and small companies was driven not by the consent decree itself but by other changes at Bell Laboratories. Historical accounts suggest that there was an exodus of important Bell researchers

[30] In Appendix C.2.7 we use different definitions for young and small companies and find that the effect is mainly driven by companies that file their first patent.

[31] We identify companies as all assignees that are never inventors. Our results are robust to defining companies as having Inc., Corp., Co. or similar abbreviations in their name.

[32] Young and small assignees are responsible for an increase of 0.014 citations (column 2). This is 70% of the total increase of 0.02 (column 1). It is also an increase of around 32% relative to what we would have expected without a consent decree. According to the estimates a Bell patent should have received 0.044 citations (0.068 is the constant, the Bell effect is -0.008, and the average decrease in citations in the post treatment period is -0.016) but did receive 0.058 citations (0.044 baseline effect + 0.014 treatment effect).

Table 3.3: The Effect of Compulsory Licensing on Subsequent Citations by Company Type and Field

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Base	Age and Size			Former Bell?		Citations to Y&S by	
	line	Young & Small	Others	Y&S Comp	No	Yes	Others	Communication
Bell	-0.4	-0.8***	0.4	-0.5**	-1.2**	0.8***	-0.6***	-0.2
	(0.5)	(0.3)	(0.5)	(0.2)	(0.5)	(0.2)	(0.2)	(0.3)
I(55-60)	-6.4***	-1.6***	-2.7***	-0.5*	-5.9***	-0.5***	-1.1***	-0.6*
	(0.6)	(0.3)	(0.8)	(0.2)	(0.6)	(0.1)	(0.2)	(0.3)
Bell x I(55-60)	2.0***	1.4***	0.6	1.0***	2.5***	-0.5**	1.1***	0.3
	(0.6)	(0.3)	(0.6)	(0.3)	(0.5)	(0.2)	(0.3)	(0.2)
Constant	18.3***	6.8***	12.4***	3.7***	17.2***	1.1***	5.2***	1.7***
	(1.2)	(0.4)	(1.0)	(0.2)	(1.2)	(0.1)	(0.7)	(0.5)
# treated	4533	4533	4533	4533	4533	4533	4533	4533
Clusters	225	225	225	225	225	225	225	225
Obs.	896556	896556	896556	896556	896556	896556	896556	896556

Notes: This table shows the results from a difference-in-differences estimation with 1949-1954 as the pre-treatment period and 1955-1960 as the treatment period. The variable Bell is an indicator variable equal to one if a patent is secured by a Bell System company before 1949 and therefore a subject to the consent decree. As control patents we use all patents that were secured in the U.S., matched by publication year, primary USPC technology class, and the number of citations up to 1949. To adjust for the different number of control patents per treatment patent in each stratum, we use the weights suggested by Iacus et al. (2009). As dependent variable, we use all citations by companies other than the filing companies in column 1. We split these citations according to the age and size of the company. In column 2, we use only citations by young and small inventors, defined as having applied for their first patent no more than ten years ago and having less than ten patents overall. In column 3, we use only the citations of inventors that are neither young nor small and in column 4 of companies that are both young and small. In columns 5 and 6, we split the citations according to whether inventors ever patented for Bell or ever were co-authors with Bell inventors. In columns 7 and 8, we split citations coming from inside or outside of the communication field. We determine whether a citing patent is inside the communication field if the technology class has the most likely application in the production of telecommunications equipment (SIC 3661), using the data of Kerr (2008). The patent data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office. All coefficients are multiplied by 100 for better readability. Standard errors are clustered on the primary three-digit USPC technology class level and *, **, *** denote statistical significance on 10%, 5% and 1% level, respectively.

around the time of the consent decree. For example, in 1953 Gordon Teal, inventor of a method to improve transistor performance, joined the then small Texas Instruments Inc. Similarly, William Shockley, one of the inventors of the transistor, left Bell in 1956 to found Shockley Semiconductors Laboratory.

To show that this is not the case, we separately look at patent citations by people who were at some point associated with Bell, but later patented for a different company, including their co-inventors, and compare with citations by all remaining unrelated inventors. In our data, there are 4,477 former Bell employees with 28,569 patents. These people have in total 12,068 co-inventors who were never active at Bell and who filed 87,148 patents in total. The results are reported in columns 5 and 6 of Table 3.3. We find a positive effect on the citations of unrelated inventors and a negative effect on the citations of related inventors. This pattern does not suggest that the increase in follow-on innovation was driven by former Bell employees. However, the results do suggest that the Bell inventors had preferential access to Bell technology prior to the consent decree and that there was a strong increase from unrelated inventors afterwards.

3.5 Compulsory Licensing did not End Foreclosure in the Market for Telecommunications Equipment

The aim of the consent decree was to end foreclosure in the market for telecommunications equipment. According to the antitrust lawsuit, Bell was closing the market to all other buyers and sellers of telecommunications equipment by using exclusive contracts between Western Electric and the Bell operating companies and by refusing to license patents to competitors. In markets outside of the telecommunications industry Bell was active only as a supplier of technology but was not an active market participant.

Market foreclosure is thought to have a negative effect on the innovation activities of the companies that are foreclosed (Baker, 2012; Wu, 2012). The argument is that foreclosed companies cannot earn profits by selling their improved products directly to consumers. The only option they have is to sell their innovations to other companies.[33] Thus, foreclosed companies have lower incentives for innovation than companies with access to a customer base.

In this section we compare the innovation effects of compulsory licensing inside and outside of the telecommunications industry to infer whether market foreclosure is harmful for innovation and whether compulsory licensing is effective in ending it. If compulsory licensing increases innovation in the same way in all industries, then any difference between the two competitive settings must be due to market foreclosure in the telecommunications industry. If market foreclosure reduces innovation as argued above and if compulsory licensing was effective in ending it, we should see a stronger increase in follow-on innovations in the telecommunications industry than in other industries. In contrast, if compulsory licensing was ineffective in ending market foreclosure, we should find a smaller effect. If market foreclosure has no effect on innovation, we should find similar effects in all industries.

To compare the innovation effects within telecommunications and outside we first need to characterize each citing patent by its closeness to the market for telecommunications equipment. To do this, we use the concordance of Kerr (2008) that gives us the probability for each USPC technology class that a patent in this technology class is used in the production of telecommunications equipment (SIC 3661). We interpret this probability as a measure of closeness to telecommunications. We then assign this probability to each citing patent according to its technology class and sum up the citations for each level of likelihood to construct a different dependent variable for each level of closeness, 26 altogether. In a last step, we repeat our main regression for each level of closeness. We can thus estimate how much the consent decree increased citations in markets that are close to the production of telecommunications equipment and in markets unrelated to it.

In Figure 3.9 we show the average treatment effects estimated with our baseline model in Equation 3.2 for different levels of closeness to the production of telecommunications equipment. We find a strong negative relation between the closeness to telecommunications and excess citations. Almost all excess citations come from patents that have nothing to do with

[33] Such a market for ideas exists only in special circumstances (Gans et al., 2002; Gans and Stern, 2003; Gans et al., 2008).

Figure 3.9: Excess Citations by Patents with Varying Likelihood of Being used in Production of Communication Equipment

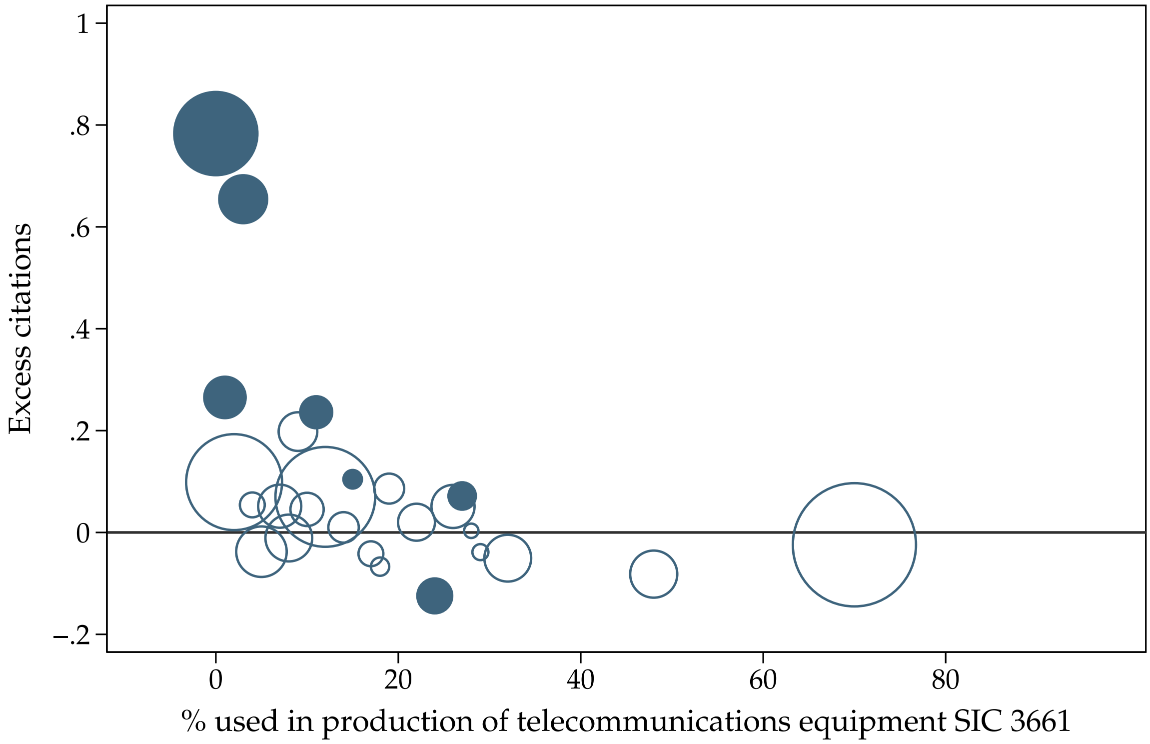

Notes: This figure shows results from a difference-in-differences estimation of the impact of the consent decree on follow-on patent citations with 1949-1954 as the pre-treatment period and 1955-1960 as the treatment period, controlling for year fixed effects. We estimate Equation 3.2 and report β_3 separately, using as dependent variables citations from patents with a different relevance for the production of telecommunication equipment (SIC 3661 - "Telephone and Telegraph Apparatus"). Relevance is measured by the likelihood that a patent is used in industry SIC 3661 using the data of Kerr (2008). The size of the circle signifies the number of Bell patents in a technology and a solid circle implies that the coefficient is significant at the 10% level. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

Figure 3.10: Excess Citations by Patents According to the Most Likely SIC Industry Classification

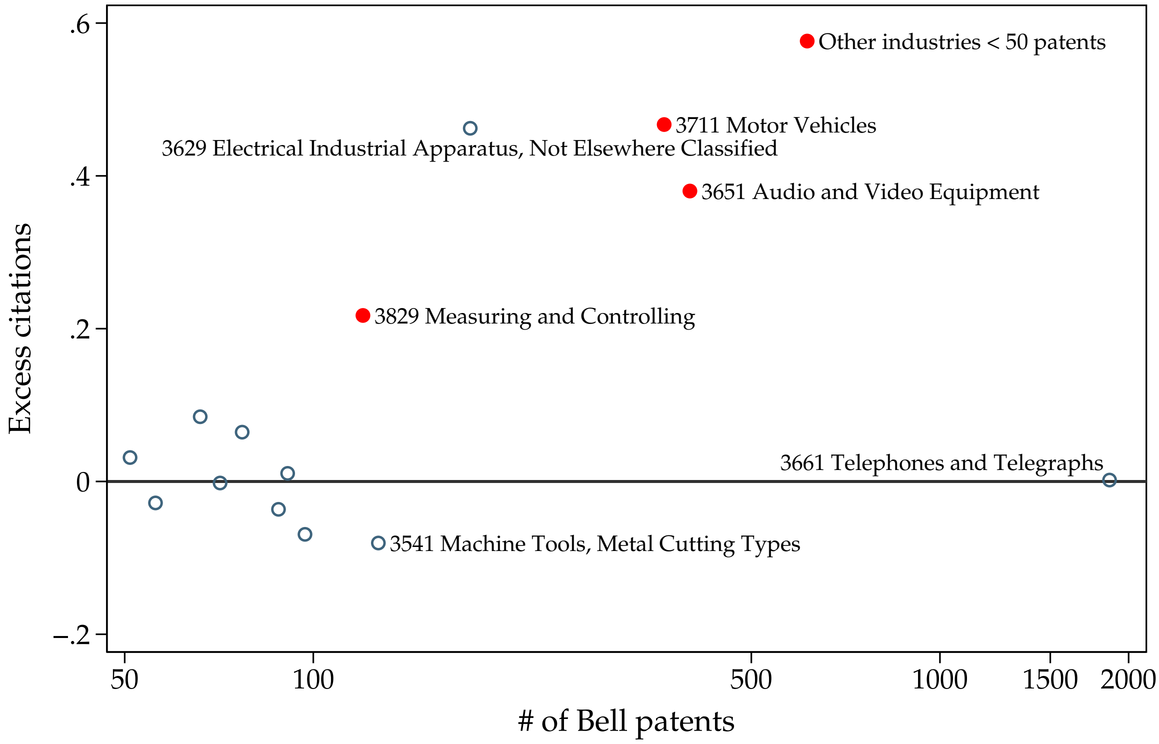

Notes: This figure shows results from a difference-in-differences estimation of the impact of the consent decree on follow-on innovation with 1949-1954 as the pre-treatment period and 1955-1960 as the treatment period, controlling for year fixed effects. As the dependent variable, we use all citations by companies other than the filing companies classified by the most likely SIC classification of the citing patent. As control patents, we use all patents that were published in the U.S. matched by publication year, primary USPC technology class, and the number of citations up to 1949. To classify a patent by its most likely industry, we use the data of Kerr (2008). We assign to each USPC class the most likely four-digit SIC industry in which it is used. A solid circle indicates that a coefficient is significant at the 10% level. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

telecommunications. We conclude from this that follow-on innovation in telecommunications was not influenced by compulsory licensing. Under the assumption that compulsory licensing affects innovation similarly in all industries this result supports the argument that market foreclosure has a chilling effect on innovation and indicates that compulsory licensing was ineffective in solving it.

Next, we use Kerr's data to assign each citing patent to the industry in which it is most likely used and repeat the baseline regression with citations from patents in different industries. The results are shown in Figure 3.10. Almost all additional citations are from patents with the most likely application outside of the industry "Telephones and Telegraphs" (SIC 3661). A large part of the effect is driven by unrelated industries such as "Measuring and Controlling," "Audio and Video Equipment" or "Motor Vehicles."[34] These results support the notion that market foreclosure is harmful for innovation and that compulsory licensing is ineffective as a remedy.[35]

Foreclosure seems to be particularly harmful for start-up innovation. In columns (7) and (8) of Table 3.3 we show that small and young companies increased their citations only outside the field of telecommunications, but not inside.[36] As a large part of the effect in the full sample was driven by small and young companies, this suggests that also start-ups react strongly to market foreclosure. In fields outside of telecommunications, compulsory licensing fostered innovation by small and young companies since Bell as the supplier of technology did not control product markets through vertical integration or via exclusive contracts.

Our results suggest that market foreclosure stifles innovation and that compulsory licensing is not sufficient to foster innovation without supporting structural remedies. This confirms the general perception at the time of the lawsuit. Both the public and antitrust officials were aware that because of Bell's persistent monopoly compulsory licensing would only help companies outside the telecommunications field. A witness in the congressional hearings put it succinctly: "while patents are made available to independent equipment manufacturers, no market for telephone equipment is supplied (...). It is rather a useless thing to be permitted to manufacture under patent if there is no market in which you can sell the product on which the patent is based." The Antitrust Subcommittee concluded that "The patent and technical information requirement have efficacy only so far as they permit independent manufacturers to avail themselves of patents in fields that are unrelated to the common carrier communication business carried on by the Bell System companies, and nothing more." On May 4, 1954, presiding Judge Stanley N. Barnes suggested that compulsory licensing policy for reasonable rates is "only good window dressing" but would do no good because Western

[34] In the Appendix C.3.1 we repeat the analysis using NBER technology subcategories to classifiy the citing patent. The results are the same.

[35] Another explanation for our null result in the telecommunications market would be that there was a lack of innovation potential in the telecommunication sector after 1956. To rule out this hypothetical possibility we compare the development of patents in the telecommunications sector. Results are reported in Figure C.9 in Appendix C.3.2. They show that the number of citations to Bell's telecommunications patents had a similar trend as patents outside of telecommunications and that the number of Bell's newly filed telecommunications patents shows no signs of abating after the consent decree.

[36] We use the most likely SIC code to determine the field of the citing patent.

Electric had already “achieved an exclusive position (...) and liberal licensing would not permit competitors to catch up” in the telecommunications business (Antitrust Subcommittee, 1959, pp. 108).

In the years after the consent decree, the Bell System faced repeated allegations of exclusionary behavior. By the 1960s and 1970s, a range of new firms were eager to enter the telecommunications market but Bell implemented measures to make it expensive or impossible (Wu, 2012). This led to a number of regulatory actions, for example forcing interconnections of Bell’s telephone system to the entering competitors MCI in 1971 which provided long distance services using microwave towers (Temin and Galambos, 1987; Gertner, 2012, p. 272). Eventually, the continued monopolization of the telecommunications market by Bell resulted in the 1974 antitrust lawsuits. The lawsuit mirrored almost scene by scene the case of 1949. Again Bell was charged with excluding competitors from the market of telecommunications equipment. Again, the Department of Defense intervened on the grounds of national defense. But the Reagan administration was not as accommodating as the Eisenhower administration had been and the Department of Justice was keen on going after Bell. The case ended with the break-up of the Bell System in 1983, opening up the market for telecommunications equipment for competition.

3.6 The Consent Decree Increased U.S. Innovation in the Long Run

The historical set-up of the Bell case gives us the opportunity to look also at the long-run innovation effects of a consent decree. In the previous section we have shown that the increase in follow-on citations is measurable for the first five years. This raises the question how lasting the impact of a large-scale intervention in patent rights really is. To answer this question we study the long-run impact of the case against Bell on the patent activities of firms patenting in the U.S. More specifically, we examine the increase in the total number of patents in a USPC technology subclass with a compulsorily licensed Bell patent relative to a subclass without. We employ the following empirical model

$$\#Patents_{s,t} = \beta_t \cdot I(Bell > 0)_s + Controls + \varepsilon_{s,t} \tag{3.3}$$

where the outcome variable is the total number of patents in a technology subclass s (Moser and Voena, 2012; Moser et al., 2014). The treatment variable equals one if there is at least one compulsorily licensed patent in the technology subclass. As controls, we use USPC class-year fixed effects.[37] Our sample consists of 235 classes with 6,276 subclasses of which 1,209 are treated.[38]

In Figure 3.11a we plot the number of excess patents for all patent classes. We leave out patents by Bell to focus on patenting of other companies. Starting in 1953, the number of patents in technology classes where Bell patents were compulsorily licensed increased relative to subclasses without Bell patents, and it continued to do so beyond 1960, when the last Bell patents affected by the consent decree expired. This suggests that the consent decree increased U.S. innovation in the long run.

To quantify the effect we next estimate the average yearly effect of the consent decree on the total number of patent applications for the time period 1949-1960. We employ the following difference-in-differences model:

$$\#Patents_{s,t} = \beta_1 \cdot I(Bell > 0)_s + \beta_2 \cdot I(Bell > 0)_s \cdot I[1955 - 1960] + Controls + \varepsilon_{s,t} \tag{3.4}$$

where $I(Bell > 0)_s$ is 1 if Bell has a patent in the subcategory s. As controls we use class-year fixed effects.

The coefficients are reported in Table 3.4. In the first five years alone, patent applications increased by 2.5 patent applications in treated classes (column 1). This is an increase of

[37] To follow the literature we use USPC technology classes here and not SIC classes.

[38] We exclude subclasses that did not have any patents at all before 1956 and we include only patent classes that contain subclasses that were treated and subclasses that were not.

Figure 3.11: Impact of the Consent Decree in the Long Run

(a) Annual Treatment Effects on the Number of Patent Applications

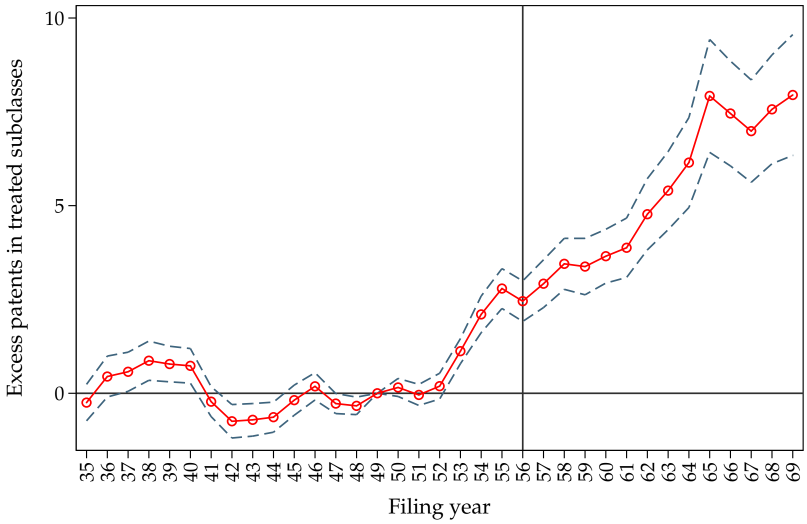

(b) Excess Patents sorted by Likelihood of Being Used in the Production of Telecommunications Equipment

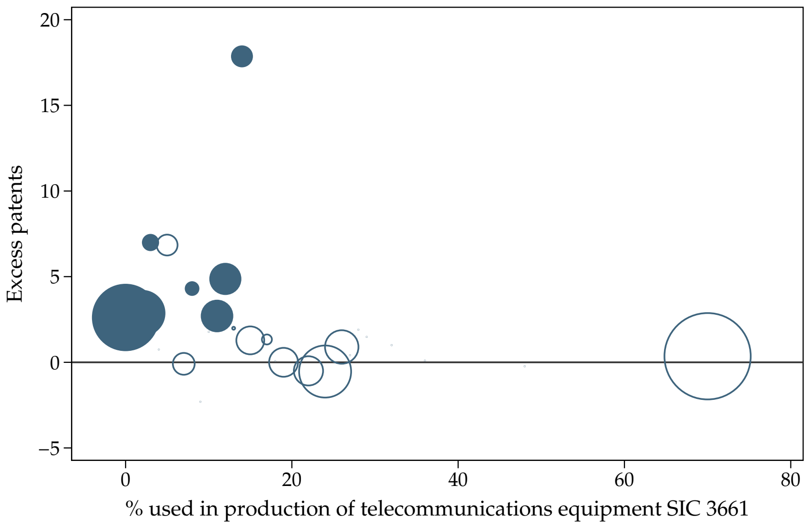

Notes: The dependent variable is the number of patent applications per (aggregated) subclass per year. A subclass is treated if it contains at least one Bell patent that was subject to compulsory licensing. Subfigure (a) shows annual treatment effects β_t estimated with Equation 3.3 for all patent classes. Standard errors are clustered at the class level. Subfigure (b) shows the average increase in the number of patents β_2 estimated with Equation 3.4 for patent classes with varying likelihood of being used in the production of telecommunications equipment. To determine the likelihood that a patent is used in industry SIC 3661 we use the data of Kerr (2008). The size of the circle signifies the number of Bell patents in a technology and a solid circle implies that the coefficient is significant on the 10% level. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

Table 3.4: Patent Applications per Subclass and Year by Company Type and Field

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Type of Company			Field			
	Baseline	Young & Small	Others	Communication	Not Comm.	Computer	Not Computer
Treated	7.5***	2.2***	5.3***	5.2***	8.2***	4.4***	7.7***
	(0.6)	(0.2)	(0.4)	(0.3)	(0.7)	(0.4)	(0.6)
Treated x I(55-60)	2.5***	1.2***	1.3***	1.4***	2.7***	3.0***	2.5***
	(0.3)	(0.1)	(0.2)	(0.5)	(0.3)	(0.8)	(0.3)
Clusters	235	235	235		211		222
Observations	75312	75312	75312	6228	69084	2268	73044

Notes: The dependent variable is the number of patent applications per subclass per year. A subclass is in the treatment group if it contains at least one Bell patent that was subject to compulsory licensing. This treatment variable is interacted with an indicator that is 1 for the period 1955-1960. The panel 1949-1960 includes subclasses that had at least one patent application between 1940 and 1949 and whose classes contain both treated and untreated subclasses. Young companies are companies with their first patent granted less than ten years ago and small companies are companies with less than 10 patents in 1949. We classify a technology as related to communication if its most likely industry is SIC 3661, the production of telecommunications equipment. The last column excludes all computer patents, defined as patents that belong to the NBER technological subcategory Computer Hardware and Software, Computer Peripherals and Information Storage (Hall et al., 2001). The patent data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office. The regressions include class-year fixed effects and standard errors are clustered at the class level, except in columns (4) and (6), where the sample is limited to 24 classes in telecommunications and 13 in computers, respectively. *, **, *** denote statistical significance on 10%, 5% and 1% level, respectively.

around 24.5%.[39] Furthermore, patent applications by new companies entering the market increased relatively more than patent applications by other companies (columns 2 and 3).[40]

The increase appears to be stronger outside of telecommunications technologies (column 4 and 5). In Figure 3.11b we plot the average treatment effects estimated with Equation 3.4 for different levels of closeness to the production of telecommunications equipment. Again the effects are weak for technologies closely related to the production of telecommunications equipment and strong for unrelated technologies. This again suggests that the fields in which Bell continued to operate experienced slower technological progress than markets where entry of start-ups was possible.

Figure 3.11a shows that the increase in patenting begins in 1953, two years before the increase in citations to Bell patents. In 1953, Bell's most important invention, the transistor, became available for licensing, spurring the creation of the computer industry. To make sure that the entire increase is not driven by this one exceptional invention, we analyze computer and non-computer patents separately and report the results in columns 6 and 7 of Table 3.4. The effect is stronger for the computer patents, but the increase in patenting is also significant without any computer patent.

Thus, overall we find that the consent decree led to a long-lasting increase in the scale of innovation mainly outside the telecommunications field. This is consistent with the theoretical argument by Acemoglu and Akcigit who build on the step-by-step innovation model of Aghion et al. (2001) to analyze the effects of compulsory licensing on innovation (Acemoglu and Akcigit, 2012). They consider the case where all current and future patents in the economy are compulsorily licensed for a positive price and identify two main effects. On the one hand, compulsory licensing helps technological laggards to catch up and brings more industries to a state of intense competition. This 'composition effect' increases innovation, because companies in industries with intense competition invest more in R&D in order to become

[39]Untreated subclasses have on average 2.17 patent applications in the pre-treatment period. In these subclasses the number of patent applications increases by 0.52 from the pre- to the post-treatment period. Using the estimate for the difference between treated and untreated classes, 7.5, in column 1 of Table 3.4, we calculate the counterfactual number of applications in treated classes in the absence of compulsory licensing which is equal to 10.19 (=7.5 + 2.17+0.52). The treatment effect is 2.5. Thus, the number of patents increased relative to the counterfactual by 24.5% (=2.5/10.19).

[40]The number of patents of young and small companies increases by 38% while the number of patents of all other companies increases by 18%.

the industry leader. On the other hands, compulsory licensing reduces the time a technology leader keeps its profitable position. This 'disincentive effect' reduces the innovation and growth in the economy.

In our case, compulsory licensing was selectively applied to one company that did not participate in any market other than the telecommunications market. This enabled many new companies to enter markets with state-of-the art technology and to compete for the industry leadership with full patent protection of future inventions intact (Holbrook et al., 2000). Thus in all industries but the telecommunications industry we measure the pure composition effect without the counteracting disincentive effect. The interpretation that the consent decree helped to open up new markets and enabled new start-ups to compete is consistent with historical accounts on the growth of electronics and computers industry in the 1950s and 1960s (Grindley and Teece, 1997).

3.7 Case Study: The Diffusion of the Transistor Technology

In this section we examine the diffusion of the transistor technology because it is a particularly insightful case study for the mechanisms illustrated in the previous sections for three reasons: First, in response to the antitrust lawsuit Bell started already in 1952 to license the transistor technology via standardized non-discriminatory licensing contracts. This creates an interesting variation in the timing of licensing. Second, transistor patents were expected to be particularly important, hence we can estimate how the amount of follow-on innovation varies with patent quality. And finally, under the impression of the antitrust lawsuit Bell was very careful not to engage in exclusionary practices with its transistor patents. Thus, in 1956 the only change for the transistor technology was that the patents were now royalty free. This allows us to examine the isolated impact of a decrease in royalties.

The transistor is arguably the most important invention of Bell Labs. As the most basic element of modern computers, the transistor has been instrumental in the creation of entire industries and its invention heralded the beginning of the information age. The invention of the transistor earned John Bardeen, Walter Brattain, and William Shockley the Nobel Prize in Physics in 1956. They filed patents in June 1948 and announced the invention on July 1 of the same year. The patents were published in 1950 and 1951. Bell, the military, and the research community at large immediately understood the importance and value of the transistor.

Due to the ongoing antitrust lawsuit, Bell's management was reluctant to draw attention to its market power by charging high prices for transistor components or for licenses (Mowery, 2011). To appease the regulator, Bell's top managers agreed to share and license the transistor device with standardized non-discriminatory licensing contracts (Gertner, 2012, p.111). In addition, Bell decided to actively promote the transistor by organizing conferences to explain the technology. In April 1952, over 100 representatives from 40 companies gathered for a nine-day Transistor Technology Symposium, including a visit to Western Electric's transistor manufacturing plant in Allentown, PA. After the conference, 30 companies decided to license the transistor technology for a non-refundable advance payment of $25,000 (~ $220,000 in today's dollars) that was credited against future royalty payments (Antitrust Subcommittee, 1958, p.2957). Royalty rates amounted to 5% of the net selling price of the transistor in 1950, which were reduced to 2% in 1953 (Antitrust Subcommittee, 1959, p. 117).

To be able to separately analyze the transistor we identify among the patents affected by the consent decree all patents related to the original transistor inventor team. There are two main transistor patents: Patent # 2,524,035 with the title "Three-Electrode Circuit Element Utilizing Semiconductive Materials" granted in 1950 to John Bardeen and Walter Brattain and Patent # 2,569,347 with the title "Circuit Element Utilizing Semiconductive Material" issued to William Shockley in 1951. To these two patents, we add all the patents of all researchers who actively worked towards the development of the transistor at Bell Labs.[41] Then we add all patents from all co-authors. We identify 329 "transistor" patents affected by the consent decree (i.e., held by Bell Labs). This sample is most likely a superset of all transistor patents. For example, it also includes patent # 2,402,662 with the title "Light Sensitive Device" granted to Russell Ohl, the original patent of the solar cell. The median publication year of the patents in the transistor subsample is 1947; and 168 of these patents are also included in our baseline sample.

[41] Researchers whom we classify to have actively contributed to the transistor at Bell Labs were in alphabetical order Bardeen, Bown, Brattain, Fletcher, Gardner Pfann, Gibney, Pearson, Morgan, Ohl, Scaff, Shockley, Sparks, Teal and Theurer (e.g. Nelson, 1962).

Figure 3.12: Annual Treatment Effects on Excess Citations of Transistor Patents

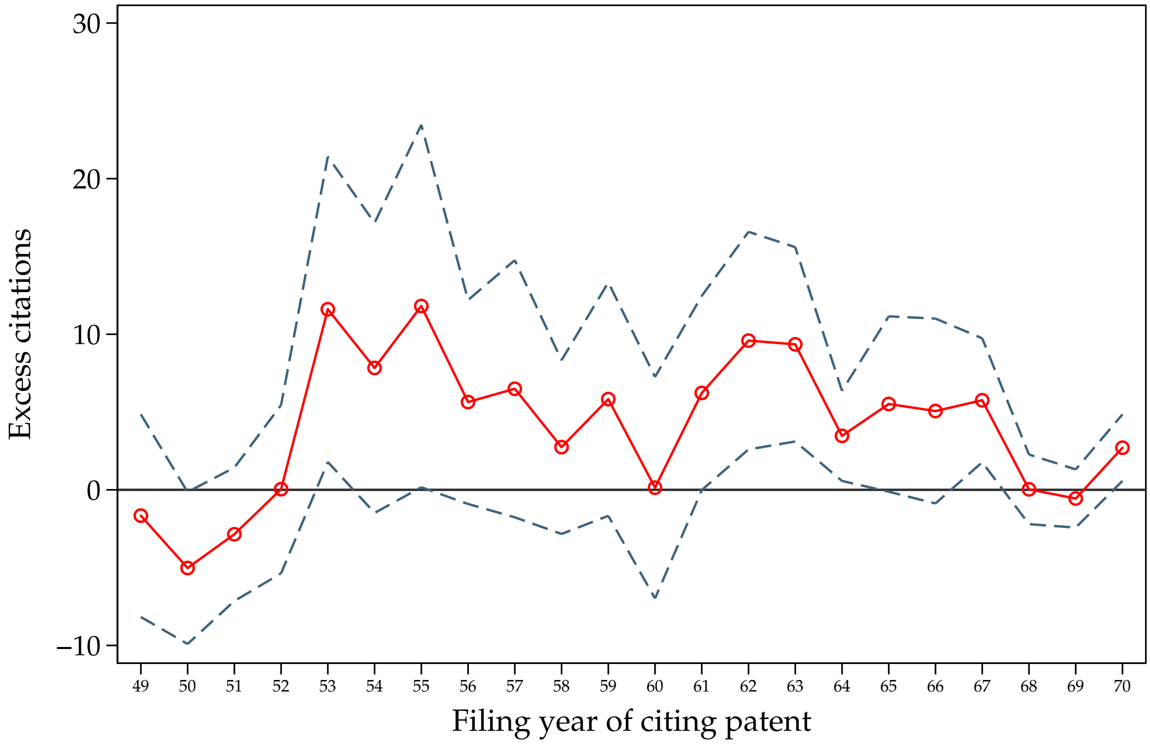

Notes: This graph shows the estimated number of yearly excess citations of transistor-related patents affected by the consent decree ("Bell patents") relative to patents with the same publication year, in the same three-digit USPC primary class and with the same number of citations up to 1951. We define Bell patents as transistor-related if they are either one of the two main transistor patents (Patent # 2,524,035 or Patent # 2,569,347) or were filed by inventors associated with these patents or their co-inventors. To arrive at these estimates, we regress the number of citations in each year on an indicator variable that is equal to one if the patent under consideration is affected by the consent decree and year fixed effects. The dashed lines represent the 90% confidence bands for the estimated coefficient. To adjust for the different number of control patents per treatment patent in each stratum, we use the weights suggested by Iacus et al. (2009). The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

To be able to repeat our regressions in this subsample of transistor patents we extend our baseline sample to patents published up to 1951. For our control group we now use patents with the same number of pre-citations up to 1951 while all other criteria stay the same.

Figure 3.12 shows the yearly excess citations of transistor patents relative to control group patents. The coefficient of 1952, which is not matched and is close to zero, speaks in favor of parallel trends. The impact of licensing is measurable starting in 1953, and lasts for at least 15 years. This suggests that standardized licensing had a positive impact on follow-on innovation. The fact that the impact does not strongly increase in 1956 when the consent

Table 3.5: The Transistor Subsample

	(1)	(2)	(3)	(4)	(5)	(6)
	Publication year <1952			Publication year<1949		
Subsample	Base-line	Transistor	No transistor	Base-line	Transis-tor	No transistor
Start treatment	1955	1953	1955	1955	1953	1955
Bell	-0.3	-1.4	-0.4	-0.4	-0.9	-0.4
	(0.3)	(1.2)	(0.3)	(0.5)	(2.1)	(0.5)
I(53/55-60)	-5.7***	-6.3**	-5.6***	-6.4***	-7.4***	-6.4***
	(0.7)	(2.7)	(0.7)	(0.6)	(2.2)	(0.6)
Bell x I(53/55-60)	1.9***	8.0**	1.8***	2.0***	4.4*	2.0***
	(0.5)	(3.7)	(0.5)	(0.6)	(2.3)	(0.6)
Constant	19.0***	23.0***	18.8***	18.3***	22.3***	18.1***
	(1.4)	(3.2)	(1.4)	(1.2)	(2.9)	(1.2)
# treated	5758	204	5554	4533	168	4365
Clusters	239	65	237	225	58	223
Obs.	1035421	64891	1021733	896556	56664	886044

Notes: This table shows the results from a difference-in-differences estimation. As the dependent variable we use all citations by companies other than the filing company. For the regression with the transistor patents, we define the treatment period as starting in 1953; for the non-transistor patents we define the treatment period as starting in 1955, as in our main regression in Equation 3.2. *Bell* is an indicator variable equal to one if a patent is published by a Bell System company before 1949 and is therefore affected by the consent decree. As control patents, we use all patents that were published in the U.S. matched by publication year, primary USPC technology class, and the number of citations. We define patents as transistor patents if they were filed by a member of the original transistor team or one of their co-authors. In the regressions for columns 1 to 3, we use all patents with a publication year before 1952 and we match all citations up to and including 1951. Correspondingly, in the regressions for columns 4 to 6 we use patents and citations up to 1949. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office. All coefficients are multiplied by 100 for better readability. Standard errors are clustered on the primary three-digit USPC technology class level and *, **, *** denote statistical significance on 10%, 5% and 1% level, respectively.

decree reduced licensing fees to zero suggests instead that the price effect of compulsory licensing had little further impact.

Table 3.5 reports the results from repeating our baseline regression in this subsample. We find that citations to the transistor patents increase by 52% (column 2). They experience around a four times higher increase in follow-on citations than other consent decree patents. The magnitude of the effect is consistent with the presumption that patents on more important inventions experience a larger increase after compulsory licensing.

Despite the large effects, the transistor patents do not drive the effect in our main sample. To rule out this possibility we analyze our original sample up to 1949 with and without transistor patents. Results are shown in columns (5) and (6). We find large but insignificant effects for the transistor sample and virtually the same effect without transistors as in the baseline regression that includes transistor patents (column 4).[42]

Transistors are the classical example of a general purpose technology that has the potential of having a large scale impact on the economy (Helpman, 1998). If it had not been for the antitrust lawsuit against Bell, odds are that Bell's licensing policy would have been less accommodating and the follow on-innovations stimulated by the transistors less dramatic than they were.

3.8 Conclusion

In this study we show that antitrust enforcement can increase innovation. The 1956 consent decree that settled the antitrust lawsuit against Bell increased innovation, mostly by small and young companies building on Bell's established technologies. We conclude that antitrust enforcement can play an important role in increasing innovation by facilitating market entry.

Several antitrust scholars have argued that antitrust enforcement should pay special attention to exclusionary practices because of their negative influence on innovation (Baker, 2012; Wu, 2012). Our study seconds this view. We show that foreclosure has a negative impact on innovation and that compulsory licensing may not be an effective remedy to end market foreclosure and to overcome its stifling effect on innovation unless accompanied by structural remedies.

Compulsory licensing is often imposed in merger cases where the market structure changes endogenously (Delrahim, 2004; Sturiale, 2011). We would expect that if the newly merged company is able to foreclose the product market, compulsory licensing is not an effective

[42] The large magnitude of the effect should not be taken at face value. The identifying assumption of this regression is that the control patents would have had the same number of citations as the transistor patents. In our regression this is true for 1953, but given the exceptional nature of the invention of the transistor, it is fair to assume that this trend might have diverged in later years. Furthermore, it is not absolutely clear from the historical records why Bell decided to license the transistor patents. If the licensing decision was taken because of the expectation of important follow-on research, our estimate might give an upper bound on the effect. For example, Jack Morton, the leader of Bell Labs' effort to produce transistors at scale, advocated the sharing of the transistor to benefit from advances made elsewhere. Source: http://www.computerhistory.org/siliconengine/bell-labs-licenses-transistor-technology/ (last accessed 9 September 2016).

remedy. More empirical studies are needed to assess whether the negative effect of market foreclosure on innovation is a first order concern for merger and acquisition cases.

We estimate the negative effects of patents on follow-on innovations by other companies, but we cannot determine how large the incentive effect of patents for the company holding the patent is. In our case, compulsory licensing does not appear to have had a strong negative effect on Bell's patenting activities. It would be surprising if this was the norm (Williams, 2015). But it is consistent with Galasso and Schankerman (2015a) who show that large companies do not reduce their innovation activity when their patents are invalidated in court, but do change the direction of their research and development activities.

We analyze a very important antitrust lawsuit from the 1950s. Using a historical setting has the advantage that we can draw on a large number of detailed historical accounts and that we can conduct a long run evaluation many years after the case. At the same time it is unclear whether the size of the effects of compulsory licensing would be similar today. Jaffe and Lerner (2011) suggest that many negative effects of the patent system discussed today are related to regulatory changes surrounding the establishment of the Court of Appeals for the Federal Circuit in 1982. The reforms led to a significant broadening and strengthening of the rights of patent holders and consequently to a surge in the number of patents granted. This makes us think that the effects of compulsory licensing might be even larger today.

Appendices

Appendix A

Appendix to Chapter 1

A.1 Additional Tables and Graphs

Figure A.1: Migration Flows, Annual Treatment Effects of Free Labor Mobility

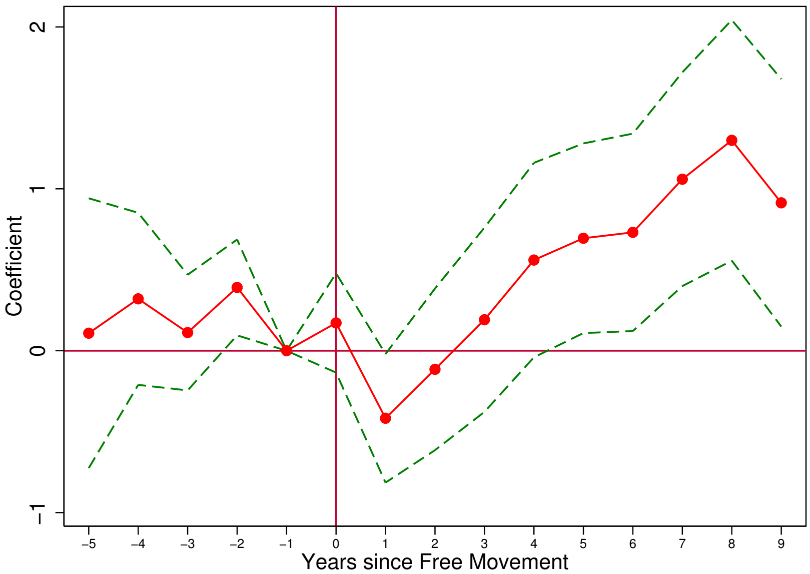

Notes: Annual treatment effects on migration around the introduction of free movement (1986-2012). The regression includes destination-year and country-pair-industry fixed effects. Standard errors are clustered at the country-pair-industry level. Source: PATSTAT, European Commission, own calculations.

Figure A.2: Inventor Mobility in Europe

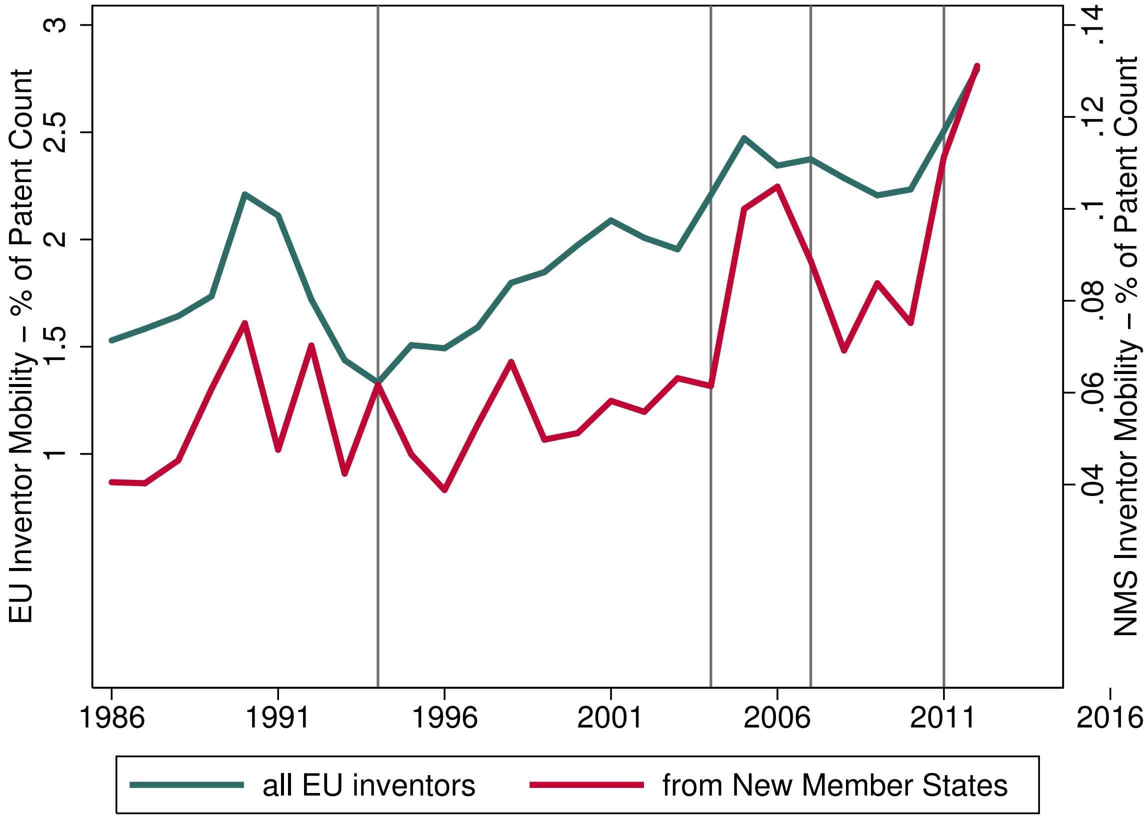

Notes: The graph shows the number of mobile inventors normalized to the total number of patent applications. We count as mobile inventor an inventor who changes his country of residence compared to the previous patent application. Thus migrants can be identified only if they have at least one patent application in each country. Source: PATSTAT.

Table A.1: Overview of the Gradual Opening of the EU15+4 Labor Markets

Country	NMS8 (2004 entry)	NMS2 (2007 entry)	Sectoral Exceptions
Austria	2011	2014	NMS8 (2007-2010), NMS2 (2007-2013): Construction, Manufacturing of Electronics and Metals, Food and beverage services (restaurant business), other sectors with labor shortages
Belgium	2009	2014	-
Denmark	2009	2009	-
Finland	2006	2007	-
France	2008	2014	NMS8 (2005-2007), NMS2 (2007-2013): Agriculture, Construction, Accommodation and food services (tourism and catering), other sectors with labor shortages
Germany	2011	2014	NMS8 (2004-2010), NMS2 (2007-2013): sectors with labor shortages
Greece	2006	2009	-
Iceland	2006	2012	-
Ireland	2004	2012	-
Italy	2006	2012	NMS8 (2004-2005): sectors with labor shortages; NMS2 (2007-2011): Agriculture, Construction, Engineering, Accommodation and food services (tourism and catering), Domestic work and care services, other sectors with labor shortages; Occupations: Managerial and professional occupations
Luxembourg	2008	2014	NMS2 (2007 - 2013): Agriculture, Viticulture, Accommodation and food services (tourism and catering)
Netherlands	2007	2014	NMS8 (2004-2006), NMS2 (2007-2013): International transport, Inland shipping, Health, Slaugtherhouse/meat-packaging, other sectors with labor shortages
Norway	2009	2012	NMS8 (2004-2008), NMS2 (2007-2011): sectors with labor shortages
Portugal	2006	2009	-
Spain	2006	2009	Reintroduction of restrictions for Romanians: 11/08/2011 - 31/12/2013
Sweden	2004	2007	-
United Kingdom	2004	2014	NMS2 (2007-2013): Agriculture, Food manufacturing

Notes: Column 2 shows the year of the labor market opening of the respective country for the NMS8 countries, column 3 shows the year of the labor market opening of the respective country for the NMS2 countries. Column 4 shows, which sectors were exempt from restrictions.
Source: European Commission.

Table A.2: Migration and Free Labor Mobility (First Stage)

	(1) EU19 and NMS all migrants	(2) NMS all migrants	(3) NMS 2004 only all migrants	(4) EU19 and NMS patent potential
L3.FM	2.352***	5.039**	19.37*	-0.563
	(0.754)	(2.320)	(10.48)	(0.645)
L4.FM	1.860***	3.271*	4.298	1.156**
	(0.630)	(1.704)	(4.065)	(0.506)
L5.FM	-0.136	-0.0996	9.662	0.350
	(0.375)	(0.418)	(19.23)	(0.292)
in EU	0.447**	-4.541*		0.261
	(0.204)	(2.526)		(0.180)
L2.Trade flow	-1.077		-74.12	-2.072**
	(1.089)		(76.99)	(0.953)
L2.FDI inflow	1.14e-05	0.000161***	0.000185***	1.45e-05**
	(2.40e-05)	(4.31e-05)	(4.76e-05)	(6.90e-06)
Observations	383	186	163	383
R-squared	0.597	0.683	0.701	0.363
Region industry FE	yes	yes	yes	yes
Year FE	yes	yes	yes	yes
F	10.20	48.57	660.0	10.11
Clusters	53	30	23	53

Notes: The regressions in this table estimate the first stage corresponding to Table 1.1 in column 1 and 4: The dependent variable is (the second lag of the natural logarithm of) emigration in a region and outflow of migrants with patenting potential, respectively. The instruments are the free movement variables for the three previous years. The regressions include controls for EU membership, trade flows and FDI inflows. The first pair of columns includes all EU and EFTA countries, the third and fourth column limit the sample to new member states and the last two columns include only the 2004 accessions. All specifications include year and region-industry fixed effects. Robust standard errors are clustered at the region-industry level. Sources: Patstat, Eurostat, CEPII
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table A.3: Patent Applications and Free Labor Mobility (Reduced Form)

	(1) EU19 and NMS Patents	(2) EU19 and NMS cit. weighted	(3) NMS Patents	(4) NMS cit. weighted	(5) NMS 2004 only Patents	(6) NMS 2004 only cit. weighted
L3.FM	1.075*	1.309*	-0.276	-0.0181	1.758	2.047
	(0.576)	(0.717)	(2.315)	(2.193)	(3.247)	(4.016)
L4.FM	1.786***	2.206***	-0.606	-0.216	-4.447	-4.335
	(0.276)	(0.386)	(0.863)	(0.805)	(3.655)	(3.624)
L5.FM	-0.177	0.0565	-0.395	-0.264	3.418	4.612
	(0.392)	(0.526)	(0.545)	(0.710)	(3.751)	(3.579)
in EU	0.167	0.278**				
	(0.107)	(0.121)				
L2.Trade flow	-1.399**	-1.456*				
	(0.662)	(0.863)				
L2.FDI inflow	3.10e-05**	4.29e-05***	1.45e-05	4.15e-05	2.49e-05	5.34e-05
	(1.26e-05)	(1.28e-05)	(2.70e-05)	(2.75e-05)	(3.15e-05)	(3.20e-05)
Observations	496	496	209	209	184	184
R-squared	0.442	0.742	0.267	0.177	0.257	0.162
Region industry FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Clusters	56	56	32	32	24	24

Notes: The dependent variables in the regressions shown in this table are the number of patent applications (columns 1, 3 and 5) and citation-weighted patent applications (columns 2, 4 and 6). More precisely, the dependent variable is the natural logarithm of 1 plus these counts. The same transformation is applied to the trade flow regressor and for FDI inflows, the percentage change from the previous year is used as regressor. The first pair of columns includes all EU and EFTA countries, columns 3 and 4 include all countries which joined the EU in 2004 and later and the last two columns only include those which joined in 2004. All specifications include year and region-industry fixed effects. Standard errors are clustered at the region-industry level. Sources: Patstat, Eurostat, CEPII

*** p<0.01, ** p<0.05, * p<0.1

Table A.4: Patent Applications and Migration in NMS10, OLS and 2SLS

	(1) OLS Patents	(2) OLS cit. weighted	(3) OLS Patents	(4) 2SLS Patents	(5) 2SLS cit. weighted	(6) 2SLS Patents
L2.Migrants	0.0924**	0.0730*		0.115	0.212	
	(0.0350)	(0.0375)		(0.156)	(0.249)	
L2.Migr.pat.potential			0.203*			0.101
			(0.112)			(0.0950)
L2.Trade flow				0.482	-0.650	0.758***
				(0.518)	(0.820)	(0.251)
L2.FDI inflow	-1.41e-05	-5.82e-06	-3.23e-07	-1.80e-05	-3.03e-05	9.41e-07
	(1.87e-05)	(1.80e-05)	(1.72e-05)	(3.34e-05)	(4.83e-05)	(1.68e-05)
Observations	163	163	163	163	163	163
Region industry FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Clusters	23	23	23	23	23	23
F				16.81	26.31	65.74

Notes: The regressions in this table estimate the relationship between the migration flow out of a country and innovation in that country. The first three columns are estimated with OLS and the last three columns use a 2SLS estimation with our instrument based on free movement legislation. The dependent variables are the number of patent applications in an industry and origin region in a year or, in columns 2 and 5, the citation-weighted patent applications (i.e., patent applications + forward citations to these patents). Patent application numbers and citation-weighted counts, number of migrants and trade flows are taken in natural logarithms. The sample includes only the 10 countries which joined the EU in 2004. All specifications include year and region-industry fixed effects. Robust standard errors are clustered at the region-industry level. Sources: PATSTAT, Eurostat, CEPII
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table A.5: Patent Applications and Migration, USPTO Patents Only, OLS and 2SLS

	(1) OLS Patents	(2) OLS cit. weighted	(3) OLS Patents	(4) 2SLS Patents	(5) 2SLS cit. weighted	(6) 2SLS Patents
L2.Migrants	0.0270	-0.0894		0.346**	0.503**	
	(0.0535)	(0.0694)		(0.171)	(0.232)	
L2.Migr.pat.potential			0.000889			0.702
			(0.0606)			(0.429)
in EU	0.0258	0.402	0.0182	0.115	0.567	-0.0508
	(0.204)	(0.305)	(0.206)	(0.238)	(0.374)	(0.222)
L2.Trade flow	1.623**	2.409**	1.740**	0.252	-0.144	2.493***
	(0.687)	(1.002)	(0.654)	(0.901)	(1.246)	(0.842)
L2.FDI inflow	1.24e-05	3.06e-05**	1.28e-05	7.25e-06	2.11e-05***	1.13e-06
	(1.02e-05)	(1.41e-05)	(9.91e-06)	(1.43e-05)	(8.16e-06)	(1.16e-05)
Observations	383	383	383	383	383	383
Region industry FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Clusters	53	53	53	53	53	53
F				32.56	273.0	26.87

Notes: The regressions in this table estimate the relationship between the migration flow out of a country and innovation in that country, counting only patents that were filed with the USPTO. The first three columns are estimated with OLS and the last three columns use a 2SLS estimation with our instrument based on free movement legislation. The dependent variables are the number of patent applications in an industry and origin region in a year or, in columns 2 and 5, the citation-weighted patent applications (i.e., patent applications + forward citations to these patents). Patent application numbers and citation-weighted counts, number of migrants and trade flows are taken in natural logarithms. The sample includes all EU members and countries in the European Free Trade Association. All specifications include year and region-industry fixed effects. Robust standard errors are clustered at the region-industry level. Sources: PATSTAT, Eurostat, CEPII
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table A.6: Convergence in Patenting Levels ($Patents_{dest}/Patents_{origin}$) and Migration, NMS only, OLS and 2SLS

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	OLS	2SLS	2SLS	2SLS
	$log(\frac{P_{diy}}{P_{oiy}})$	$log(\frac{P_{diy}}{P_{oiy}})$	$log(\frac{P_{diy}}{P_{oiy}})$	$log(\frac{P_{diy}}{P_{oiy}})$	$log(\frac{P_{diy}}{P_{oiy}})$	$log(\frac{P_{diy}}{P_{oiy}})$
	Patents	cit. weighted	Patents	Patents	cit. weighted	Patents
L2.Migrants	0.0289	0.0349		-0.254*	-0.259*	
	(0.0229)	(0.0285)		(0.141)	(0.151)	
L2.Migr.pat.potential			0.118			-1.809
			(0.0747)			(2.810)
Patents, origin	-1.080***	-1.052***	-1.081***	-1.083***	-1.055***	-1.065***
	(0.109)	(0.120)	(0.109)	(0.110)	(0.122)	(0.119)
Patents, dest	1.078***	1.128***	1.080***	1.071***	1.120***	1.044***
	(0.0713)	(0.0910)	(0.0717)	(0.0765)	(0.0962)	(0.0890)
Within EU	0.0435	-0.0278	0.0431	0.0572	-0.0136	0.0724
	(0.0529)	(0.0584)	(0.0533)	(0.0555)	(0.0612)	(0.0693)
GDP_d/GDP_o	-0.444	-0.0942	-0.446	-0.480	-0.132	-0.471
	(0.359)	(0.413)	(0.360)	(0.372)	(0.430)	(0.365)
L3.Trade flow	-0.0394	0.0516	-0.0355	0.00323	0.0959	-0.0277
	(0.0623)	(0.0792)	(0.0616)	(0.0662)	(0.0842)	(0.0633)
L3.FDI flow	0.00139	0.00150	0.00108	0.000924	0.00101	0.00544
	(0.00662)	(0.00662)	(0.00663)	(0.00750)	(0.00741)	(0.0112)
Observations	2,763	2,763	2,763	2,681	2,681	2,681
R-squared	0.499	0.565	0.499	0.458	0.535	0.406
Origin-dest-industry FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Clusters	559	559	559	477	477	477
F				90.89	137.0	81.10

Notes: The dependent variable is the natural logarithm of $Patents_{dest}/Patents_{origin}$. Number of patents (in origin and destination countries), number of migrants, FDI, and trade flows are in natural logarithms. The sample includes country-industry pairs, where origins are NMS and destinations - EU19 countries. All specifications include year and origin-destination-industry fixed effects. Robust standard errors are clustered at the origin-destination-industry level. Sources: PATSTAT, Eurostat, CEPII
*** p<0.01, ** p<0.05, * p<0.1

Table A.7: Convergence in Patenting Levels ($Patents_{dest}/Patents_{origin}$) and Free Labor Mobility (Reduced Form)

	(1) EU19 and NMS Patents	(2) EU19 and NMS cit. weighted	(3) NMS only Patents	(4) NMS only cit. weighted	(5) EU19 and NMS (all) Patents	(6) NMS only (all) Patents
L3.FM	-0.0135	0.00186	-0.0150	-0.00166	0.0179	-0.0131
	(0.0368)	(0.0434)	(0.0412)	(0.0496)	(0.0122)	(0.0130)
L4.FM	-0.0631	-0.0573	-0.0534	-0.0554	-0.0403***	-0.0337**
	(0.0440)	(0.0469)	(0.0505)	(0.0554)	(0.0133)	(0.0146)
L5.FM	-0.0256	-0.0393	-0.0267	-0.0283	-0.0166	-0.00647
	(0.0419)	(0.0449)	(0.0495)	(0.0534)	(0.0127)	(0.0137)
Patents, origin	-1.242***	-1.407***	-1.094***	-1.067***	-0.640***	-0.618***
	(0.0797)	(0.0873)	(0.111)	(0.122)	(0.0113)	(0.0114)
Patents, dest	1.051***	1.090***	1.062***	1.112***	0.800***	0.813***
	(0.0725)	(0.0921)	(0.0729)	(0.0929)	(0.0216)	(0.0218)
Within EU	-0.00662	-0.100*	0.0241	-0.0442	-0.0781***	-0.0527***
	(0.0494)	(0.0549)	(0.0553)	(0.0620)	(0.0141)	(0.0149)
GDP_d/GDP_o	0.00771	0.555	-0.251	0.0737	0.183***	0.175***
	(0.331)	(0.391)	(0.384)	(0.451)	(0.0393)	(0.0402)
L3.Trade flow	-0.0450	0.00903	-0.0127	0.0766	-0.0499***	-0.0341***
	(0.0629)	(0.0810)	(0.0629)	(0.0807)	(0.00866)	(0.00878)
L3.FDI flow	0.00170	0.000342	0.00259	0.00241	-0.0140***	-0.0112***
	(0.00656)	(0.00665)	(0.00651)	(0.00659)	(0.00416)	(0.00418)
Observations	2,946	2,946	2,763	2,763	71,496	66,504
R-squared	0.487	0.552	0.500	0.565	0.217	0.225
Origin-dest-ind FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Clusters	582	582	559	559	5688	5304

Notes: The dependent variable is the natural logarithm of $Patents_{dest}/Patents_{origin}$. Number of patents (in origin and destination countries), number of migrants, FDI, and trade flows are in natural logarithms. All specifications include year and origin-destination-industry fixed effects. Robust standard errors are clustered at the origin-destination-industry level. Specifications 1-4 show the reduced form regressions for the sample used in the OLS/2SLS estimations (i.e., the sub-sample for which migration data are available), specifications 5-6 show estimates for the full sample of country-industry pairs in 2000-2012. Sources: PATSTAT, Eurostat, CEPII
*** p<0.01, ** p<0.05, * p<0.1

Table A.8: Citations to Destination Industries, NMS only, OLS and 2SLS

	(1) OLS	(2) OLS	(3) OLS	(4) 2SLS	(5) 2SLS	(6) 2SLS
L2.Migrants	0.00255	0.00895		0.427*	0.436*	
	(0.0281)	(0.0282)		(0.222)	(0.224)	
L2.Migr.pat.potential			0.124			5.695
			(0.133)			(5.457)
Patents, origin		0.124***	0.124***		0.146***	0.158**
		(0.0332)	(0.0331)		(0.0369)	(0.0622)
L3.Patents, dest		0.0118	0.0121		0.0183	0.0317
		(0.0224)	(0.0224)		(0.0248)	(0.0332)
Within EU		-0.00869	-0.00991		-0.0280	-0.0827
		(0.0608)	(0.0609)		(0.0637)	(0.0991)
L3.Trade flow		-0.0575	-0.0566		-0.122	-0.0791
		(0.0722)	(0.0724)		(0.0837)	(0.0895)
L3.FDI flow		0.00342	0.00306		0.00393	-0.0129
		(0.0122)	(0.0122)		(0.0127)	(0.0268)
Observations	2,763	2,763	2,763	2,681	2,681	2,681
R-squared	0.083	0.087	0.088			
Origin-dest-industry FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Clusters	559	559	559	477	477	477
F				11.64	8.404	6.418

Notes: The dependent variable is the number of citations from a region and industry to another country in a year. Citation counts, number of migrants, total number of patent applications in origin and destination industries, FDI and trade flows are taken in natural logarithms. The sample is limited to new EU member states. All specifications include year and origin-destination-industry level. Robust standard errors are clustered at the origin-destination-industry level. Sources: PATSTAT, Eurostat, CEPII
*** p<0.01, ** p<0.05, * p<0.1

Table A.9: Citations to Destination Industries and Free Labor Mobility (Reduced Form)

	(1)	(2)	(3)	(4)
			EU19 and NMS	NMS only
	EU19 and NMS	NMS only	(all)	(all)
L3.FM	0.00662	0.0785	0.0400**	0.0670**
	(0.0349)	(0.0516)	(0.0163)	(0.0306)
L4.FM	0.0734	0.0856	0.0431**	0.0903**
	(0.0451)	(0.0603)	(0.0182)	(0.0392)
L5.FM	0.0480	0.0753	0.0255	0.0406
	(0.0470)	(0.0559)	(0.0169)	(0.0337)
Patents, origin	0.138***	0.134***	0.0591***	0.0974***
	(0.0238)	(0.0308)	(0.00785)	(0.0119)
L3.Patents, dest	0.0478***	0.000519	0.0258***	-0.0130*
	(0.0137)	(0.0192)	(0.00541)	(0.00687)
Within EU	0.0154	0.163***	0.00274	0.180***
	(0.0361)	(0.0553)	(0.0144)	(0.0247)
L3.Trade flow	-0.152***	-0.0732	-0.0627***	0.0372**
	(0.0352)	(0.0596)	(0.00955)	(0.0144)
L3.FDI flow	-0.000418	0.0114	0.0257***	0.0235***
	(0.00520)	(0.0113)	(0.00357)	(0.00639)
Observations	7,279	3,498	29,604	11,851
R-squared	0.174	0.133	0.099	0.110
Origin-dest-industry FE	yes	yes	yes	yes
Year FE	yes	yes	yes	yes
Clusters	1322	592	2304	912

Notes: The dependent variable is the number of citations from a region and industry to another country in a year. Citation counts, number of migrants, total number of patent applications in origin and destination industries, FDI and trade flows are taken in natural logarithms. All specifications include year and origin-destination-industry level. Robust standard errors are clustered at the origin-destination-industry level. Columns 1 and 2 show the reduced form regressions for the sample used in the OLS/2SLS estimations (i.e., the sub-sample for which migration data are available), columns 3 and 4 show estimates for the full sample of country-industry pairs in 2000-2012. Sources: PATSTAT, Eurostat, CEPII

*** p<0.01, ** p<0.05, * p<0.1

Table A.10: Citations to Destination Industries, USPTO Patents Only, OLS and 2SLS

	(1) OLS	(2) OLS	(3) OLS	(4) 2SLS	(5) 2SLS	(6) 2SLS
L2.Migrants	0.0476**	0.0313*		0.679***	0.288	
	(0.0188)	(0.0178)		(0.197)	(0.184)	
L2.Migr.pat.potential			0.0512			0.745
			(0.0485)			(1.703)
Patents, origin		0.193***	0.194***		0.186***	0.195***
		(0.0221)	(0.0221)		(0.0227)	(0.0227)
L3.Patents, dest		0.0545***	0.0542***		0.0542***	0.0491**
		(0.0147)	(0.0147)		(0.0150)	(0.0195)
Within EU		0.00444	0.00724		-0.00468	0.0300
		(0.0332)	(0.0332)		(0.0343)	(0.0667)
L3.Trade flow		0.0797*	0.0858**		0.0294	0.0853**
		(0.0416)	(0.0417)		(0.0552)	(0.0418)
L3.FDI flow		-0.00960*	-0.0102*		-0.00738	-0.0152
		(0.00526)	(0.00527)		(0.00568)	(0.0134)
Observations	7,299	7,287	7,287	7,136	7,124	7,124
R-squared	0.132	0.150	0.149			
Origin-dest-industry FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Clusters	1322	1320	1320	1159	1157	1157
F				44.41	35.32	34.64

Notes: The dependent variable is the number of citations from a region and industry to another country in a year. Citation counts, number of migrants, total number of patent applications in origin and destination industries, FDI and trade flows are taken in natural logarithms. The sample is limited to citations among U.S. patents. All specifications include year and origin-destination-industry level. Robust standard errors are clustered at the origin-destination-industry level. Sources: PATSTAT, Eurostat, CEPII

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table A.11: Only Citations Added by the Applicant

	(1) OLS	(2) OLS	(3) OLS	(4) 2SLS	(5) 2SLS	(6) 2SLS
L2.Migrants	0.0234	0.0258		0.489***	0.336*	
	(0.0176)	(0.0172)		(0.170)	(0.192)	
L2.Migr.pat.potential			0.0759**			1.239
			(0.0342)			(1.915)
Patents, origin		0.149***	0.150***		0.139***	0.150***
		(0.0220)	(0.0220)		(0.0231)	(0.0233)
L3.Patents, dest		0.0253	0.0247		0.0249	0.0162
		(0.0161)	(0.0161)		(0.0165)	(0.0221)
Within EU		-0.0992***	-0.0958***		-0.110***	-0.0574
		(0.0335)	(0.0334)		(0.0354)	(0.0712)
L3.Trade flow		-0.0194	-0.0143		-0.0805	-0.0153
		(0.0369)	(0.0369)		(0.0542)	(0.0373)
L3.FDI flow		0.00811	0.00735		0.0108*	-0.000953
		(0.00506)	(0.00508)		(0.00554)	(0.0154)
Observations	7,299	7,287	7,287	7,136	7,124	7,124
R-squared	0.070	0.080	0.080			
Origin-dest-industry FE	yes	yes	yes	yes	yes	yes
Year FE	yes	yes	yes	yes	yes	yes
Clusters	1322	1320	1320	1159	1157	1157
F				22.90	20.46	18.95

Notes: The dependent variable is the number of citations from a region and industry to another country in a year. Citation counts, number of migrants, total number of patent applications in origin and destination industries, FDI and trade flows are taken in natural logarithms. The sample is limited to citations which have been added by the applicant according to PATSTAT. Robust standard errors are clustered at the origin-destination-industry level. Source: Eurostat and PATSTAT.
*** p<0.01, ** p<0.05, * p<0.1

Appendix B

Appendix to Chapter 2

B.1 Proof of Proposition 1

First note that $\frac{d}{d\Delta}\left(\frac{I_L}{I_L+I_S}\right) < 0$ is equivalent to $\frac{d}{d\Delta}\left(\frac{n_S}{n_L}\right) > 0$, since

$$\frac{I_L}{I_L + I_S} = \frac{2\mu_0 n_L}{2\mu_0 n_L + 2\mu_0 n_S} = \frac{n_L}{n_L + n_S} = \frac{1}{1 + n_S/n_L}.$$

Thus the proof only requires showing that $\frac{d}{d\Delta}\left(\frac{n_L}{n_S}\right) < 0$. Plugging the values for profits into equations 2.3 and 2.4 gives optimal investments.

$$n_L = \frac{1}{8} + \frac{\Delta}{4}$$
$$n_S = \frac{\Delta}{2}$$

It can now be seen that the ratio $\frac{n_L}{n_S}$ is a decreasing function of Δ.

$$\frac{n_L}{n_S} = \frac{1}{4\Delta} + \frac{1}{2}$$
$$\frac{\partial\left(\frac{n_L}{n_S}\right)}{\partial\Delta} = \frac{-1}{4\Delta^2} < 0$$

B.2 Proof of Proposition 2

The expected absolute amounts of small and large innovation can be expressed in the following way, using the values for R&D investments in the leveled state from the proof of Proposition 1.

$$I_S = 2\mu_0 n_S = \mu_0 \Delta$$
$$I_L = 2\mu_0 n_L = \mu_0 \left(\frac{1}{4} + \frac{\Delta}{2}\right)$$

Next the steady state share of industries being in the leveled state (or equivalently, the probability of any particular industry being in that state) needs to be found. Using Equation 2.6 and the fact that these probabilities have to sum to one $\mu_0 + \mu_F = 1$ and recalling from above that $n_F = \frac{1}{8} - \frac{\Delta}{4}$, this can be written as:

$$\mu_0 = \frac{n_F + h}{2(n_L + n_S) + n_F + h} = \frac{1 + 8h - 2\Delta}{3 + 8h + 10\Delta}.$$

This leads to the following expressions for I_L and I_S

$$I_L = \frac{\frac{1}{4} + 2h + 4h\Delta - \Delta^2}{3 + 8h + 10\Delta}$$
$$I_S = \frac{(1 + 8h)\Delta - 2\Delta^2}{3 + 8h + 10\Delta}$$

The derivatives of these functions with respect to the level of competition are as follows.

$$\frac{dI_L}{d\Delta} = \frac{-\frac{5}{2} - 8h + 32h^2 - (6 + 16h)\Delta - 10\Delta^2}{[3 + 8h + 10\Delta]^2} \tag{B.1}$$
$$\frac{dI_S}{d\Delta} = \frac{3 + 32h + 64h^2 - (12 + 32h)\Delta - 20\Delta^2}{[3 + 8h + 10\Delta]^2} \tag{B.2}$$

That the absolute innovation levels are indeed concave functions of competition for the relevant values of h can be shown in the following way: Let $x\prime$ denote the derivative of a variable x with respect to Δ and $i \in \{L, S\}$.

$$\begin{aligned} I_i &= 2\mu_0 n_i \\ I_i\prime &= 2(\mu_0\prime n_i + \mu_0 n_i\prime) \\ I_i\prime\prime &= 2(\mu_0\prime\prime n_i + 2\mu_0\prime n_i\prime + \mu_0 n_i\prime\prime) \end{aligned}$$

The derivatives in these expressions are:

$$\begin{aligned} \mu_0\prime &= \frac{-2[3+8h+10\Delta] - 10(1+8h-2\Delta)}{[3+8h+10\Delta]^2} = \frac{-(16+96h)}{[3+8h+10\Delta]^2} < 0 \\ \mu_0\prime\prime &= 20(16+96h)[3+8h+10\Delta]^{-3} > 0 \\ n_L\prime &= \frac{1}{4} > 0 \\ n_S\prime &= \frac{1}{2} > 0 \\ n_L\prime\prime &= n_S\prime\prime = 0 \end{aligned}$$

Note that $\mu_0\prime = -\frac{1}{20}[3+8h+10\Delta]\mu_0\prime\prime$, such that $I_i\prime\prime$ can be expressed as:

$$I_i\prime\prime = 2(\mu_0\prime\prime n_i + 2\mu_0\prime n_i\prime) = 2\mu_0\prime\prime(n_i - \frac{1}{10}[3+8h+10\Delta]n_i\prime).$$

It remains to be shown when the term in brackets is negative for the two innovation technologies, such that I_i is concave. Consider first the large step, for which the following inequality must then hold.

$$\begin{aligned} n_i &< \frac{1}{10}[3+8h+10\Delta]n_i\prime \\ n_L = \frac{1}{8} + \frac{\Delta}{4} &< \frac{1}{10}[3+8h+10\Delta]\frac{1}{4} \\ 40(\frac{1}{8} + \frac{\Delta}{4}) &< 3+8h+10\Delta \\ 5 &< 3+8h \\ \frac{1}{4} &< h \end{aligned}$$

The corresponding inequality for the small step always holds.

$$n_S = \frac{\Delta}{2} < \frac{1}{10}[3 + 8h + 10\Delta]\frac{1}{2}$$
$$10\Delta < 3 + 8h + 10\Delta$$
$$0 < 3 + 8h$$

Thus I_S is always concave and I_L is concave if $\frac{1}{4} < h$. If innovation rates are non-negative at $\Delta = 0$, for an inverted-U shape to emerge they must be decreasing at $\Delta = \frac{1}{2}$. With the equations above, we can see that the condition $\frac{dI_S}{d\Delta}|_{\Delta=\frac{1}{2}} < 0$ implies $h < \frac{1}{4}$. If this is not satisfied, the level of small innovation increases with competition for any initial level of competition.

$\frac{dI_L}{d\Delta}$ shows that for I_L to be increasing at $\Delta = 0$ the condition $h > \frac{1+\sqrt{6}}{8} \approx 0.43$ must be satisfied, such that the condition for concavity is satisfied as well ($\frac{1}{4} < h$). The condition $\frac{dI_L}{d\Delta}|_{\Delta=\frac{1}{2}} < 0$ implies $h < \frac{1+\sqrt{5}}{4} \approx 0.81$. (Hence, if I_L is not concave, it is decreasing for the entire range of competition intensities.) If the help parameter h satisfies both of these conditions I_L is inverted-U shaped.

The level of total innovation, $I = I_L + I_S$, has similar properties. It is a concave function of Δ, since $I'' = I_L'' + I_S'' < 0$ holds for any h. Summing up the derivatives $\frac{dI_L}{d\Delta} + \frac{dI_S}{d\Delta}$ from above, we can see that at $\Delta = 0$ the slope will always be positive, i.e., for low levels of competition, increasing the intensity of competition always increases the total level of innovation. I is inverted-U shaped if at $\Delta = \frac{1}{2}$ this sum is negative. This is the case if $h < \sqrt{\frac{1}{6}} \approx 0.41$. If this condition is not satisfied, overall innovation is increasing in competition for any initial level.

These results are summarized in the proposition.

B.3 Proof of Proposition 3

From the proof of Proposition 2, we know that the I and I_S are concave and that if I_L is not always decreasing in competition, it is concave as well. Thus the maxima are reached where the first derivatives are zero or at the boundary of the range of possible values of Δ, i.e., at $\Delta = \frac{1}{2}$ or $\Delta = 0$, following from proposition 2 as well.

If the maxima are not at the boundaries, then the curves are inverted-U shaped and have interior maxima. If I_L has an interior maximum, it will be where $\frac{dI_L}{d\Delta} = 0$. At this value for Δ the numerator of Equation B.1 is equal to zero.

$$-\frac{5}{2} - 8h + 32h^2 + (-6 - 16h)\Delta - 10\Delta^2 = 0$$
$$10\Delta^2 + (6 + 16h)\Delta + \frac{5}{2} + 8h - 32h^2 = 0$$

Thus, if the maxima for the I_L, I_S and I curves are interior, they are (the latter two are derived in the same way using Equation B.2 and the sum of equations B.1 and B.2, respectively):

$$\Delta_L^* = \frac{-(3+8h) + \sqrt{(3+8h)^2 - 25 - 80h + 320h^2}}{10}$$
$$\Delta_S^* = \frac{-(3+8h) + \sqrt{(3+8h)^2 + 15 + 160h + 320h^2}}{10}$$
$$\Delta^* = \frac{-(3+8h) + \sqrt{(3+8h)^2 + \frac{5}{3} + 80h + 320h^2}}{10}$$

From the above equations for the maxima it can be seen that the number of large innovations in the economy is maximized at a lower Δ than that of small innovations. Using the properties derived above, one can also see that this must hold in the following way.

$$I_L = I\frac{I_L}{I_L + I_S}$$
$$I_L\prime = I\prime\frac{I_L}{I_L + I_S} + I\frac{d}{d\Delta}\left(\frac{I_L}{I_L + I_S}\right) < 0$$

At an interior maximum of total innovation I, the first order condition requires that the first summand be zero. For the second summand, we know that I is positive and, from Proposition 1 above, that the summand must then be negative. Thus, decreasing Δ marginally from the Δ^* that maximizes I increases I_L. The reverse holds for I_S, i.e., the number of small innovations increases moving to higher levels of competition from an interior Δ^*.

$$I_S = I\frac{I_S}{I_L + I_S}$$

$$I_S\prime = \underbrace{I\prime \frac{I_S}{I_L + I_S}}_{=0} + \underbrace{I \frac{d}{d\Delta}\left(\frac{I_S}{I_L + I_S}\right)}_{>0} > 0$$

B.4 Model with Fully Incorporated Product Market Competition

The model in the main part of the chapter makes various simplifications in order to allow for an analytical solution. This section verifies with numerical solutions that these are not necessary for the results described in the propositions for the main model and discusses how the two models are related. In particular, the model in this appendix modifies the following features of the previous model:

- Firms look ahead infinitely and discount the future appropriately rather than choosing R&D investments only to maximize the expected profit in the next period.
- The model is in continuous time, such that innovations naturally do not occur at exactly the same time and there is no need to ensure that a certain upper limit on investments is not exceeded in equilibrium. Recall that in the previous model, only one or zero innovations could occur per period and thus, the sum of the probabilities of innovations for the two types and two firms could not exceed 1. The limit on equilibrium investment was introduced implicitly through the cost functions.
- Costs for different types of innovation are likely to differ, e.g. because high expected return comes with high risk of failure. As mentioned above, different costs for different types of innovations affect innovation probabilities in equilibrium and can ensure, for example, that large innovations occur less frequently than small ones.
- A large step may lead to a large, but finite, cost reduction, such that the competitor remains active in the market.
- Finally, the profit levels at different technological positions all follow directly from the product market competition model here, such that a follower's position and profit differ depending on the step size by which the competitor leads, as depicted in Figure B.1.

Figure B.1: Possible Positions on the Technology Ladder

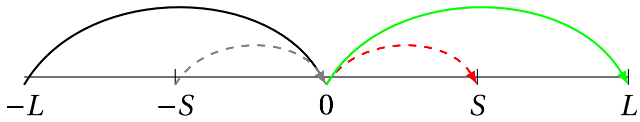

Notes: Possible positions on the technology ladder: Firms may be equally advanced (0), lead by a small step (S) or by a large step (L), or lag behind by a small ($-S$) or a large step ($-L$).

This implies that investments in catch-up innovation will also differ between these two positions. Substitutability α is the measure of competition in this model, replacing Δ.

Duopolists compete in the product market, producing output at their current costs. In addition to the output decision, firms invest in two different R&D technologies, which affect the probability of generating cost-saving innovations. The framework builds on existing models of innovation (Aghion et al., 2001, 2005), but departs in that firms choose how much to invest in the two types of innovation technologies. Neck-and-neck firms determine their R&D strategy with two continuous variables, denoting the probability of a successful large or small step, respectively. Again, a cost function for each innovation type determines the costs associated with a certain success probability.

To justify the central assumptions regarding profits in the R&D model in the main part, I start from a product market competition model following Aghion et al. (2001). In each industry sector there are two firms offering differentiated products and competing in prices. A unit mass of consumers maximizes the utility function $u = \int_{s=0}^{1} \ln((q_{sA}^{\alpha} + q_{sB}^{\alpha})^{\frac{1}{\alpha}})ds$, with $0 < \alpha \leq 1$, which integrates over all sectors s. Because of the natural logarithm in the utility function, consumer spending in each sector is the same and it is convenient to normalize this amount to 1, i.e., prices are expressed such that in all sectors $p_{sA}q_{sA} + p_{sB}q_{sB} = 1$. This leads to the following demand for firm $i \in \{A, B\}$:[1]

$$q_i = \frac{p_i^{\frac{1}{\alpha-1}}}{p_A^{\frac{\alpha}{\alpha-1}} + p_B^{\frac{\alpha}{\alpha-1}}} \tag{B.3}$$

[1] Industry indices have been omitted from these equations, since all sectors are alike.

Firm i produces q_i with constant marginal costs c_i, such that its first order condition is[2]

$$\alpha - \alpha(p_i - c_i)q_i - \frac{c_i}{p_i} = 0. \tag{B.4}$$

Since both equalities hold for each firm, this constitutes a system of four equations, which defines prices and quantities for both firms at the product market stage. Costs depend on previous research outcomes and may be considered fixed at the production stage. Thus, with the above system, the endogenous variables p_i and q_i determine profits π_i for the firms as well (disregarding any R&D costs at this point).

$$\pi_i = (p_i - c_i)q_i \tag{B.5}$$

Aghion et al. (2001) derive some properties of the implicitly defined profit $\pi_i(\frac{c_i}{c_j}, \alpha)$, where c_i denotes the costs of producing one unit of output for firm $i, j \in \{A, B\}$ and $i \neq j$. Profit depends only on relative costs and the intensity of competition, measured by the substitutability parameter α. Note that this does not imply that industry profit is constant: When one firm innovates and is a step ahead of its competitor, industry profit increases despite constant consumer spending. The innovator can benefit from decreased costs because of the competitor's higher costs.

Profits in the symmetric equilibrium and the leveled state, i.e., when the duopolists are at the same technological position, can be found as follows. (Indices are omitted due to symmetry and equal costs in this industry state.)

$$q = \frac{1}{2p} \tag{B.6}$$

$$\alpha - \alpha(p - c)q - \frac{c}{p} = 0. \tag{B.7}$$

$$\Rightarrow q = \frac{\alpha}{2(2-\alpha)c}, \quad p = \frac{2-\alpha}{\alpha}c \tag{B.8}$$

[2] The firm chooses a price to maximize $\pi_i = (p_i - c_i)q_i$ with the demand function above.

This implies that each firm's profit in the leveled state is $\pi_0 = \frac{1-\alpha}{2-\alpha}$ (where the index 0 denotes the relative technological position of the firms – in this case the technological distance to the competitor is zero).[3]

At this point, it becomes necessary to formalize the notion of small and large steps. The motivation for this study is more the potential decline in very large innovations than the nature of the smaller steps that are observed. Thus, large innovations[4] can be modeled as those that result in a monopoly for one firm, which corresponds to the extreme case in the model where the innovating firm's costs drop to zero (while the competitor's costs remain positive). Small innovations are those that result in reduced but positive costs for the innovator. Furthermore, I restrict the small innovations considered in the model to those that are large enough to increase an innovator's profits compared to profits when monopolizing half of the market (which is the case when $\alpha \to 0$). The assumption of a duopoly makes the requirement appear restrictive, in that it requires "small" innovations to be rather large. While the precise distinction corresponding to the main model would indeed actually be between infinite steps and (sufficiently large but) finite steps, the requirement on the size of the inventive step for small innovations appears much more likely to be satisfied when we think of a global market with numerous small national markets. Then even a small cost advantage that can be turned into profit on the entire quantity of a product sold globally (Bertrand competition and complete substitutability through open borders and cheap transportation) may well be worth more than the monopoly profit in a single small economy. Hence the model should be interpreted qualitatively, illustrating mechanisms and describing the directions of effects rather than quantifying their sizes. If, as in the numerical solution below, large steps lead to very small, but positive, marginal costs, the qualitative results remain the same as in the extreme case of zero marginal costs for large innovations.

The intuition for the property that small innovations become on average relatively more attractive compared to large ones (with the aforementioned restriction) as substitutability (α) increases from a value near 0 to 1 is then easy to see in Figure B.2. The other axis denotes

[3] Cf. eq. 8, Aghion et al. (2001). Unlike in their model, here the number of steps between competitors is not explicitly counted as there are only four relative positions a firm can be in, as explained below.

[4] These innovations resemble "drastic innovations", following Tirole (1988). If marginal costs drop to zero, the innovator extracts the entire (fixed) consumer spending of 1 as profit and is able to offer an infinite amount of the good in return. Thus, the price per unit of the good goes to 0 and is below any strictly positive costs of the competitor, such that the competitor is no longer a constraint on the innovator's price setting.

Figure B.2: Product Market Profits

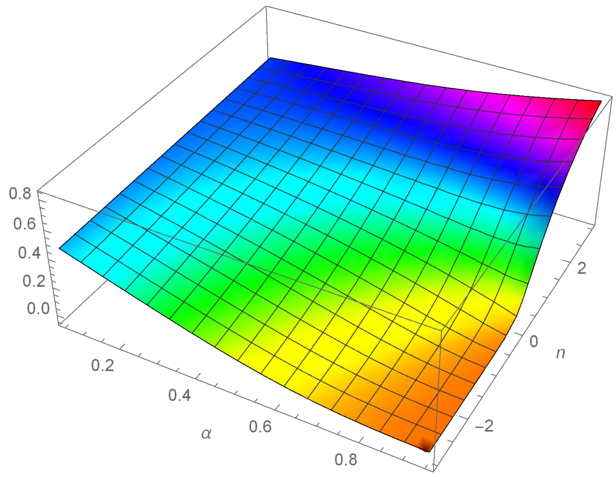

Notes: The figure shows firm profits depending on the intensity of competition α and the distance on the technology ladder (where the cost ratio is 2^{-n} and n can be understood as a measure of the firm's technological lead over its competitor).

the cost ratio, i.e. the size of the inventive step. E.g. at $n = 1$, the cost ratio is $\frac{1}{2}$, meaning that the leading firm produces at half the costs of the other firm.[5] The profits of neck-and-neck firms decrease from $\frac{1}{2}$ to 0 over this range of α. This decrease makes both types of innovations more attractive, as they have been defined to lead to non-decreasing post-innovation profits. The crucial difference between the two types is that the large innovation always leads to a monopoly in the entire market and thus a fixed profit of 1 (the maximum possible profit, since consumer spending is assumed constant). Post-innovation profits after a small innovation are close to $\frac{1}{2}$ for low levels of α, thus the difference between pre-innovation and post-innovation profits is also close to zero. For $\alpha = 1$, however, the firms are in Bertrand competition with perfect substitutability between their products. Again by normalization of market size to 1, the innovator's profit is equal to the cost difference in this case, which by the assumption on the minimum size of small steps is larger than $\frac{1}{2}$. Thus incentives to invest in both types of innovation increase as neck-and-neck profits decrease.[6]

[5] The n here only refers to the step size and not to R&D investments.

[6] As Figure B.2 illustrates, the laggard's profit after a small innovation is also decreasing in the intensity of competition. After a large innovation, the laggard's profit is always zero (as a result of the assumption that a large innovation is "drastic"). Note that if large steps were less extreme and led to small but still positive costs, the qualitative results of this study would go through, as demonstrated with the parameterization for the numerical solution below. The difference is just that increasing post-innovation profits would increase the incentives to invest in large R&D for very low levels of α. As post-innovation costs approach zero (or step size approaches infinity), post-innovation profit will approach its maximum faster and faster as α increases starting from zero.

As the difference between pre-innovation and post-innovation profits for neck-and-neck firms is increasing in α for the entire range from 0 to 1, an increase in competition will always lead to more innovation of both types in this industry state.

The system of equations includes firms' profits in the product market under the costs brought about by the outcome of R&D projects. The numerical solutions below show that, when the product market competition model determines all the profit levels, the economy exhibits a decreasing share of large innovations with competition under reasonable parameterizations.[7] The maximum of the absolute amount of large innovations also occurs at lower levels of competition than that of either total or small innovations in this model. To allow for an analytical derivation of these results, however, in the model in the main part only the essential feature of the product market competition model is maintained by introducing the competition parameter Δ (replacing α), which directly affects profits. In the main model, a laggard's profit is defined to be $\pi_F = 0$, while the large step leader gets $\pi_L = 1$, both independently of competition. The competition parameter Δ captures the crucial feature of the product market model that increasing substitutability increases the spread between the profits of leveled firms and a firm leading by a small step. It can be defined via the product market model, which gives the profit in the leveled state π_0, such that $\Delta \equiv \frac{1}{2} - \pi_0$. The profit after a small innovation is set to $\pi_S = \frac{1}{2} + \Delta$. Recall from above that $\pi_0 = \frac{1-\alpha}{2-\alpha}$, such that Δ is an increasing function of α, with $\Delta = 0$ if $\alpha = 0$ and $\Delta = \frac{1}{2}$ if $\alpha = 1$. Therefore Δ increases from 0 to $\frac{1}{2}$ as competition intensifies.[8] In the following, these simplications are not necessary, as the model is studied only numerically.

A neck-and-neck firm's present value is the result of the following maximization:[9]

[7] To be precise, this is the case if competition is sufficiently intense. The minimum intensity for increasing competition to decrease the share of large innovations goes to zero as the large step size approaches infinity.

[8] The definition of Δ above implies that $\Delta = \frac{1}{2} - \frac{1-\alpha}{2-\alpha}$. Note that the values for π_F and π_L correspond to the values in the product market competition model for a large step which leads to zero marginal costs and any $\alpha > 1$. At $\alpha = 1$, $\Delta = \frac{1}{2}$ and the simplified profits $\pi_F = 0$ and $\pi_S = 1$ are approached by the product market model as the (finite) cost reduction brought about by the small step increases.

[9] The interest rate r is another exogenous parameter here and corresponds to the time preference of households, as in Aghion et al. (2001).

$$\begin{aligned}V_0 = \max_{\{n_L, n_S\}} [& \left(\pi_0 - \beta_L \frac{n_L^2}{2} - \beta_S \frac{n_S^2}{2}\right) dt \\ & + e^{-rdt} \big(n_L dt V_L + n_S dt V_S + \overline{n}_L dt V_{-L} + \overline{n}_S dt V_{-S} \\ & + (1 - dt(n_L + n_S + \overline{n}_L + \overline{n}_S)) V_0 \big)].\end{aligned} \tag{B.9}$$

π_0 is the profit in the product market for a neck-and-neck firm minus the total R&D costs. As in the main model, depending on R&D outcomes the firm may transition to a different state. Future values are discounted by e^{-rdt}. The chance of making a large innovation depends on n_L and, if successful, the firm becomes a leader by a large step, such that its present value will be V_L. Analogously, the transition to small step leadership follows in the equation. The competitor's corresponding investments, $\overline{n}_L$ and $\overline{n}_S$, are taken as given. The future present value may be V_{-L} or V_{-S} if the competitor had a large or a small innovation, respectively, and if no innovation occurs the present value is V_0 again.

As in the main model, a lagging firm only has to choose the level of catch-up R&D (h_{-L} and h_{-S} are the respective help parameters). However, there are two different investment levels, n_{-L} and n_{-S}, now for large and small step lags behind the competitor, respectively. In this optimization the present values are:

$$\begin{aligned}V_{-L} = \max_{n_{-L}} [& \left(\pi_{-L} - \beta_{-L} \frac{n_{-L}^2}{2}\right) dt \\ & + e^{-rdt} ((n_{-L} + h_{-L}) dt V_0 + (1 - dt(n_{-L} + h_{-L})) V_{-L})]\end{aligned} \tag{B.10}$$

$$\begin{aligned}V_{-S} = \max_{n_{-S}} [& \left(\pi_{-S} - \beta_{-S} \frac{n_{-S}^2}{2}\right) dt \\ & + e^{-rdt} ((n_{-S} + h_{-S}) dt V_0 + (1 - dt(n_{-S} + h_{-S})) V_{-S})]\end{aligned} \tag{B.11}$$

As before, a firm leading by a large or a small step, has no R&D decisions to take and no R&D costs, such that the present values do not involve a maximization.

$$V_L = \pi_L dt + e^{-rdt}\left((\overline{n}_{-L} + h_{-L})dtV_0 + (1 - (\overline{n}_{-L} + h_{-L})dt)V_L\right) \tag{B.12}$$

$$V_S = \pi_S dt + e^{-rdt}\left((\overline{n}_{-S} + h_{-S})dtV_0 + (1 - (\overline{n}_{-S} + h_{-S})dt)V_S\right) \tag{B.13}$$

Taking advantage of the fact that $e^{-rdt} \approx 1 - rdt$ as dt becomes small (cf. Aghion et al., 2001), the following Bellman equations can be derived (and terms involving $(dt)^2$ can be ignored). To briefly exemplify the derivation, consider the last equation (for V_S).

$$\begin{aligned}
V_S &= \pi_S dt + (1 - rdt)\left((\overline{n}_{-S} + h_{-S})dtV_0 + (1 - (\overline{n}_{-S} + h_{-S})dt)V_S\right) \\
V_S &= \pi_S dt + (\overline{n}_{-S} + h_{-S})dtV_0 + (1 - (\overline{n}_{-S} + h_{-S})dt)V_S - rdtV_S \\
rdtV_S &= \pi_S dt + (\overline{n}_{-S} + h_{-S})dtV_0 - (\overline{n}_{-S} + h_{-S})dtV_S \\
rV_S &= \pi_S + (\overline{n}_{-S} + h_{-S})(V_0 - V_S)
\end{aligned}$$

Applying the same transformation to the other value functions leads to the following system of Bellman equations:

$$rV_L = \pi_L + (V_0 - V_L)(n_{-L} + h_{-L}) \tag{B.14}$$

$$rV_S = \pi_S + (V_0 - V_S)(n_{-S} + h_{-S}) \tag{B.15}$$

$$\begin{aligned}
rV_0 = \pi_0 - \beta_L \frac{n_L^2}{2} - \beta_S \frac{n_S^2}{2} + (V_L - V_0)n_L + (V_S - V_0)n_S \\
+ \overline{n}_S(V_{-S} - V_0) + \overline{n}_L(V_{-L} - V_0)
\end{aligned} \tag{B.16}$$

$$rV_{-S} = \pi_{-S} - \beta_{-S}\frac{n_{-S}^2}{2} + (V_0 - V_{-S})(n_{-S} + h_{-S}) \tag{B.17}$$

$$rV_{-L} = \pi_{-L} - \beta_{-L}\frac{n_{-L}^2}{2} + (V_0 - V_{-L})(n_{-L} + h_{-L}) \tag{B.18}$$

Firms' maximization of present values (or, equivalently, annuity values) leads to the first order conditions below, which give the optimal investment levels.

$$n_L = (V_L - V_0)/\beta_L \tag{B.19}$$

$$n_S = (V_S - V_0)/\beta_S \tag{B.20}$$

$$n_{-S} = (V_0 - V_{-S})/\beta_{-S} \tag{B.21}$$

$$n_{-L} = (V_0 - V_{-L})/\beta_{-L} \tag{B.22}$$

In the steady state the shares of industries in the three states (leveled or one firm leading with a technological distance of either a large step or small step) have to sum to one: $\mu_0 + \mu_S + \mu_L = 1$. Again, inflows and outflows have to be be the same for each state.

$$\mu_0(2n_L + 2n_S) = (\mu_S + \mu_L)(n_F + h) \tag{B.23}$$

$$\mu_S(n_F + h) = 2\mu_0 n_S \tag{B.24}$$

$$\mu_L(n_F + h) = 2\mu_0 n_L \tag{B.25}$$

The average innovation rates for large and small steps are defined as in the main model (in equations 2.7 and 2.8). Unlike in the main model of the study, however, the model described in this section will be studied only numerically and all the profits in the Bellman equations follow directly from the product market competition model and the associated cost ratios.

Figures B.3 to B.5 illustrate the model with a parameterization that leads to very similar results to the simplified model in the main part of the chapter. The costs of innovations differ by type. Increasing the probability of a large innovation at the technological frontier (in the leveled state) is most expensive with the cost parameter $\beta_L = 8$ and a small innovation is half as expensive $\beta_S = 4$. Catch-up innovation from the large lag is associated with the cost parameter $\beta_{-L} = 2$ and from a small lag the parameter is $\beta_{-S} = 1$. Both help parameters h_{-L} and h_{-S} are equal to $\frac{1}{10}$ and constant marginal costs of production are equal to[10] $c = 1$, except after a small innovation, which reduces costs to $c_S = \frac{1}{100}$, or after a large innovation that leads to finite, but very small marginal costs of $c_L = \frac{1}{10000}$.

[10] More precisely, the stated costs are the *relative* costs of the innovating firm compared to the competitor's costs without innovation.

Figure B.3: Profits at Different Positions

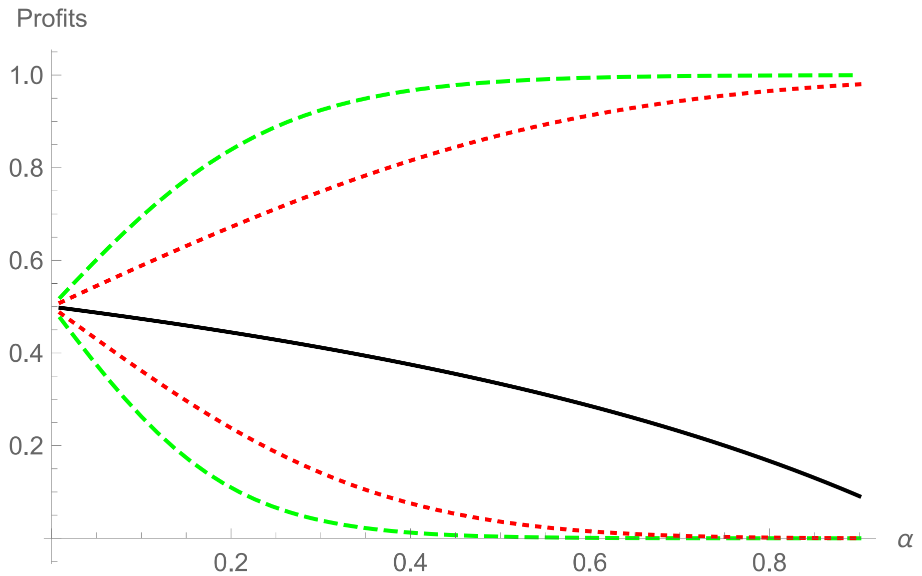

Notes: This graph shows the profits for leveled firms (black) as well as small step (upper dotted red line) and large step leaders (upper dashed green) and followers lagging with a small (lower dotted red) and large step (lower dashed green). ($r = \frac{1}{10}$, $\beta_L = 8$, $\beta_S = 4$, $\beta_{-S} = 1$, $\beta_{-L} = 2$, $h_{-L} = \frac{1}{10}$, $h_{-S} = \frac{1}{10}$, $c = 1$, $c_S = \frac{1}{100}$, $c_L = \frac{1}{10000}$)

Figure B.3 shows how profits change with competition (i.e., substitutability α) for the different firm positions. The figure shows that the mechanism from the simplified model remains: Both R&D investments become more attractive for firms in the leveled state as pre-innovation profits (π_0, black line) decrease. However, large step leader's profit (π_L, upper green line) increases earlier and flattens sooner than the profit of a small step leader (π_S, upper red line), as we move to higher values for α. This leads to the familiar shift towards large step innovation, illustrated in Figure B.4. Note that the share would not be monotonously decreasing for all parameterization. Numerical solutions suggest that if there is an initial increase, it happens at lower levels of competition the greater the large cost reduction.

Figure B.5 confirms that innovation rates for this parameterization are similar to the simplified model as well. The example shows that the three curves may continue to be inverted-U shaped.

The model in this section is considerably more complicated and numerical results are shown for one possible parameterization to demonstrate that the relationships in the simplified model may arise in a model without these abstractions.

Figure B.4: Share of Large Innovations

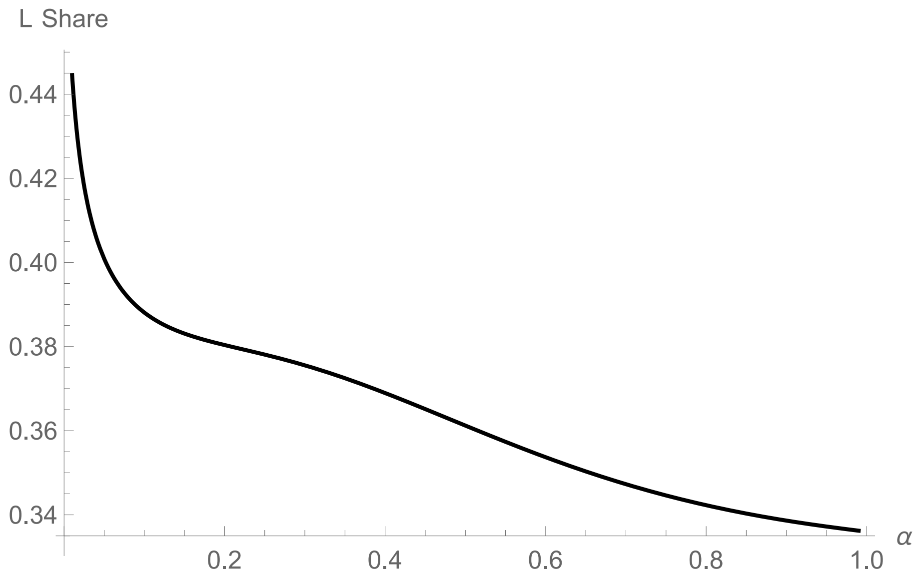

Notes: This graph shows the share of large innovations in all innovations at the technological frontier ($\frac{I_L}{I_L+I_S}$). ($r = \frac{1}{10}$, $\beta_L = 8$, $\beta_S = 4$, $\beta_{-S} = 1$, $\beta_{-L} = 2$, $h_{-L} = \frac{1}{10}$, $h_{-S} = \frac{1}{10}$, $c = 1$, $c_S = \frac{1}{100}$, $c_L = \frac{1}{10000}$)

Figure B.5: Innovation Rates

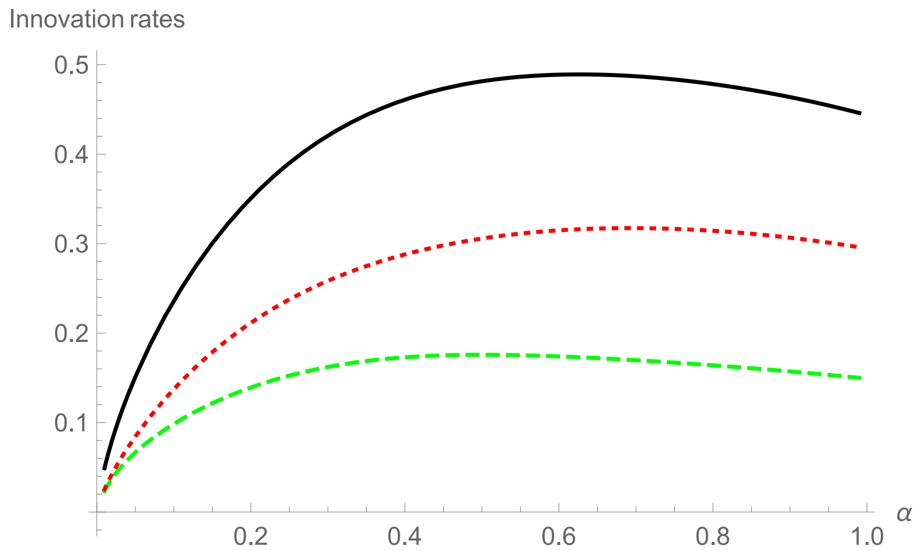

Notes: This graph shows the innovation rates for large (dashed green) and small (dotted red) innovation and their sum (black). ($r = \frac{1}{10}$, $\beta_L = 8$, $\beta_S = 4$, $\beta_{-S} = 1$, $\beta_{-L} = 2$, $h_{-L} = \frac{1}{10}$, $h_{-S} = \frac{1}{10}$, $c = 1$, $c_S = \frac{1}{100}$, $c_L = \frac{1}{10000}$)

B.5 Additional Empirical Results

Figure B.6: Kernel Density

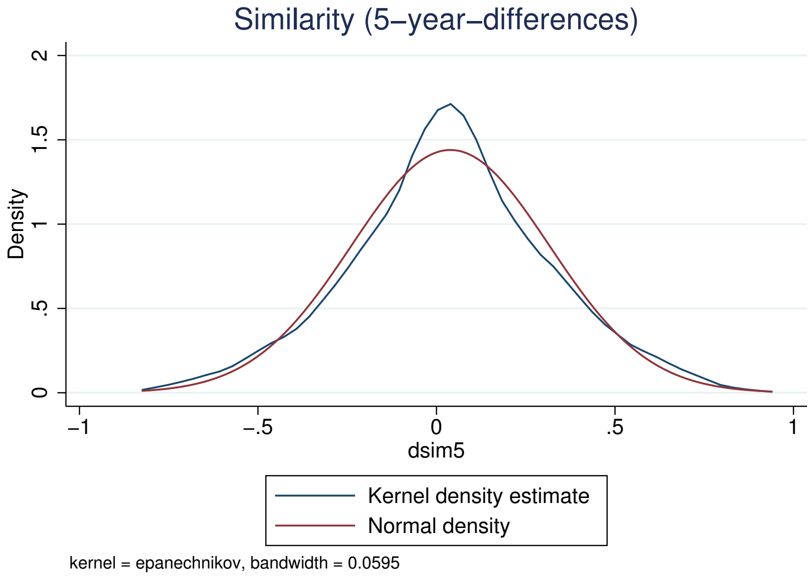

Table B.1: Similarity Regressions Controlling for Change in Patenting

	(1)	(2)	(3)	(4)	(5)
	Δ sim(comp.)	Δ sim(SIC4)	Δ sim(SIC3)	Δ sim(SIC2)	dsim5all
	b/se	b/se	b/se	b/se	b/se
Δ ShareImChina	1.486***	0.337	0.316	0.154	0.158
	(0.541)	(0.330)	(0.340)	(0.242)	(0.243)
$\Delta \ln(\text{patents}+1)$	-0.040***	-0.015*	-0.018**	-0.013*	-0.014*
	(0.015)	(0.009)	(0.008)	(0.007)	(0.008)
Year FE	yes	yes	yes	yes	yes
SIC2 FE	no	no	no	no	no
Companies	298	421	425	428	428
SIC4	114	135	139	142	142
SIC3	67	74	75	78	78
SIC2	17	19	19	19	19
N	707	860	865	868	868

Notes: This table adds a control variable for the change in patenting (more precisely, the five-year difference in ln(number of patents + 1)) to the regressions from Table 2.2. The dependent variable is the five-year difference in the similarity measure (based on SVD with 100 components). In column 1, the comparison group consists of the firm's own past patents. In columns 2 to 4, the comparison group contains patents of all firms in the same industry, where the industry is defined based on the Standard Industrial Classification (SIC) increasingly broadly (using four, three and two-digit SIC categories, respectively). In column 5 the comparison group includes patents of firms from all other industries as well. The model includes year fixed effects (as in Equation 2.10) and is estimated with two-stage least squares using Bloom et al.'s (2016) initial conditions variable to instrument for the increase in Chinese import competition. The sample period is from 2001 to 2005. Standard errors are clustered at the four-digit SIC level. The data are from PATSTAT, Bloom et al. (2016) and Peruzzi et al. (2014).

*** p<0.01, ** p<0.05, * p<0.1

Table B.2: Descriptive Statistics (Used in Regression with Similarity to Own Patents)

	mean	sd	min	p10	p25	p50	p75	p90	max
Δ sim(comp.)	0.03	0.27	-0.76	-0.30	-0.13	0.04	0.19	0.38	0.81
Δ sim(SIC4)	0.01	0.17	-0.51	-0.22	-0.10	0.01	0.13	0.22	0.66
Δ sim(SIC3)	0.00	0.16	-0.60	-0.19	-0.10	0.01	0.11	0.19	0.52
Δ sim(SIC2)	0.01	0.15	-0.46	-0.17	-0.08	0.01	0.09	0.19	0.50
Δ sim(all)	0.00	0.14	-0.49	-0.18	-0.08	0.01	0.09	0.18	0.48
Δ ShareImChina	0.02	0.04	-0.06	0.00	0.00	0.01	0.03	0.06	0.54
$\text{ShareImChina}_{jt-6}$									
$\cdot\Delta\text{ShareImChina}_t$	0.15	0.13	0.00	0.05	0.08	0.11	0.18	0.25	0.75
Observations	707								

Notes: The descriptives in this table include the observations of the estimation based on the comparison to a firm's own previous patents (see Table 2.2, column 1). This sample is the smallest one of the five columns, as the existence of own patents in the comparison period implies that the industry in aggregate has patents as well. This sample includes five-year differences for the years from 2001 to 2005. The data are from PATSTAT, Bloom et al. (2016) and Peruzzi et al. (2014).

Table B.3: Distribution of Firms over Countries

Country	Percent
GB	24.5
DE	22.5
IT	17.4
FR	16.8
SE	8.7
FI	4.0
ES	2.3
AT	2.0
DK	1.7
Total	100.0
N	298

Notes: This table shows the percentage of firms in each country for the sample used in the regressions based on similarity to a firm's own patents in Table 2.2 (column 1). The notes to Table 2.2 regarding the sample apply.

Table B.4: Distribution of Firms over Industries

SIC2	Description	Percent
35	Industrial and Commercial Machinery and Computer Equipment	24.2
28	Chemicals and Allied Products	15.1
34	Fabricated Metal Products	10.7
36	Electronic & Other Electrical Equipment & Components	10.4
38	Measuring, Photographic, Medical, & Optical Goods, & Clocks	8.7
37	Transportation Equipment	7.7
30	Rubber and Miscellaneous Plastic Products	4.4
39	Miscellaneous Manufacturing Industries	4.4
32	Stone, Clay, Glass, and Concrete Products	4.0
33	Primary Metal Industries	2.7
20	Food and Kindred Products	2.0
25	Furniture and Fixtures	2.0
26	Paper and Allied Products	1.3
24	Lumber and Wood Products, Except Furniture	0.7
27	Printing, Publishing and Allied Industries	0.7
31	Leather and Leather Products	0.7
22	Textile Mill Products	0.3
Total		100.0
N		298

Notes: This table shows the percentage of firms in each two-digit SIC category for the sample used in the regressions based on similarity to a firm's own patents in Table 2.2 (column 1). The notes to Table 2.2 regarding the sample apply. Descriptions are added from `http://www.dnb.com/content/dam/english/economic-and-industry-insight/sic_2_digit_codes.xls`.

B.5.1 Regressions Using Different Samples of Patents

Table B.5: Similarity (USPTO Patents Only)

	(1)	(2)	(3)	(4)	(5)
	Δ sim(comp.)	Δ sim(SIC4)	Δ sim(SIC3)	Δ sim(SIC2)	Δ sim(all)
	b/se	b/se	b/se	b/se	b/se
Δ ShareImChina	1.642**	-0.237	-0.390	-0.443	-0.224
	(0.686)	(0.430)	(0.469)	(0.432)	(0.346)
Year FE	yes	yes	yes	yes	yes
Companies	127	187	191	195	195
SIC4	69	84	88	92	92
SIC3	45	52	53	57	57
SIC2	14	17	17	17	17
N	270	341	346	350	350

Notes: This table replicates Table 2.2 with a restricted sample: Only patent applications filed with the USPTO are included. Apart from that, the notes to Table 2.2 apply.

Table B.6: Similarity with All Patents (Including Chinese and Other Non-English Patents)

	(1)	(2)	(3)	(4)	(5)
	Δ sim(comp.)	Δ sim(SIC4)	Δ sim(SIC3)	Δ sim(SIC2)	Δ sim(all)
	b/se	b/se	b/se	b/se	b/se
Δ ShareImChina	1.192**	0.167	0.135	0.052	0.038
	(0.498)	(0.309)	(0.317)	(0.240)	(0.242)
Year FE	yes	yes	yes	yes	yes
Companies	315	457	462	464	464
SIC4	116	136	141	143	143
SIC3	68	74	76	78	78
SIC2	18	19	19	19	19
N	764	937	943	945	945

Notes: This table replicates Table 2.2 with an extended sample, which in addition includes patents filed in China, as well as any patents for which no English abstract is available. Apart from that, the notes to Table 2.2 apply.

B.5.2 Regressions Weighting Firms Equally

Probability weights are used to weight firms equally here. I.e., individual observations for a firm, whose five-year differences in similarity are observed in two years, receive half the weight of a firm which is only once in the regression's sample.

Table B.7: Similarity Regressions Weighting Firms Equally

	(1)	(2)	(3)	(4)	(5)
	Δ sim(comp.)	Δ sim(SIC4)	Δ sim(SIC3)	Δ sim(SIC2)	Δ sim(all)
	b/se	b/se	b/se	b/se	b/se
Δ ShareImChina	1.121*	0.406	0.211	0.012	0.156
	(0.670)	(0.329)	(0.267)	(0.251)	(0.239)
Year FE	yes	yes	yes	yes	yes
Companies	298	421	425	428	428
SIC4	114	135	139	142	142
SIC3	67	74	75	78	78
SIC2	17	19	19	19	19
N	707	860	865	868	868

Notes: This table presents modified regressions from Table 2.2, where each firm in the sample receives equal weight. E.g. an observation of a long difference for a firm for which two observations are included in the sample receives half the weight of an observation for a firm for which only one long difference is observed. Apart from that, the notes to Table 2.2 apply.

Table B.8: First-Stage Regressions Weighting Firms Equally

	(1)	(2)	(3)	(4)	(5)
	Δ ShareImChina	Δ ShareImChina	Δ ShareImChina	Δ ShareImChina	Δ ShareImChina
	b/se	b/se	b/se	b/se	b/se
$\text{ShareImChina}_{jt-6}$					
$\cdot\Delta\text{ShareImChina}_t$	0.112***	0.149***	0.149***	0.149***	0.149***
	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
Year FE	yes	yes	yes	yes	yes
Obs.	707	860	865	868	868
R^2	0.11	0.19	0.19	0.18	0.18
F	15.77	24.76	25.27	25.50	25.50

Notes: This table shows the first-stage regressions belonging to the 2SLS regressions in Table B.7.

The number of observations per firm included in the regressions is shown in the Table B.9 (the sample is the same as in Table 2.2 in the main part). In the regression in the first column of Table B.7, for 40% of included firms only one long difference is observed. For 20% of firms differences are observed in two year, while the percentages of firms with three, four and five observations are above 10%. The additional observations in the four other columns of the table do not alter this distribution much.

Table B.9: Number of Observations per Firm

n	(1) comp. Percent	(2) SIC4 Percent	(3) SIC3 Percent	(4) SIC2 Percent	(5) all Percent
1	40.3	53.0	53.2	53.5	53.5
2	19.8	17.6	17.6	17.5	17.5
3	15.1	10.9	10.8	10.7	10.7
4	12.1	9.3	9.2	9.1	9.1
5	12.8	9.3	9.2	9.1	9.1
Total	100.0	100.0	100.0	100.0	100.0
N	298	421	425	428	428

Notes: This table shows the distribution of the number of observations per firm included in the regressions Table B.7 (which is the same as for Table 2.2). The columns correspond to the columns in these tables, i.e., the first column includes the sample in the regressions with the similarity to a company's own previous patents, in the second to fourth column the comparison group is the four, three and two-digit SIC category of the firm and in the fifth column all patents are included in the comparison group.

B.5.3 First-Stage and OLS Regressions

Table B.10: First Stage of Similarity Regressions

	(1) Δ ShareImChina b/se	(2) Δ ShareImChina b/se	(3) Δ ShareImChina b/se	(4) Δ ShareImChina b/se	(5) Δ ShareImChina b/se
$\text{ShareImChina}_{jt-6}$ $\cdot\Delta\text{ShareImChina}_t$	0.127***	0.145***	0.146***	0.146***	0.146***
	(0.022)	(0.023)	(0.023)	(0.023)	(0.023)
Year FE	yes	yes	yes	yes	yes
Obs.	707	860	865	868	868
R^2	0.14	0.17	0.17	0.17	0.17
F	12.54	16.37	16.53	16.70	16.70

Notes: This table shows the first-stage regressions belonging to the main 2SLS regressions in Table 2.2.

Table B.11: Similarity (OLS)

	(1)	(2)	(3)	(4)	(5)
	Δ sim(comp.)	Δ sim(SIC4)	Δ sim(SIC3)	Δ sim(SIC2)	Δ sim(all)
	b/se	b/se	b/se	b/se	b/se
Δ ShareImChina	0.480*	0.337**	0.108	0.113	0.038
	(0.250)	(0.161)	(0.113)	(0.099)	(0.100)
Year FE	yes	yes	yes	yes	yes
Companies	298	421	425	428	428
SIC4	114	135	139	142	142
SIC3	67	74	75	78	78
SIC2	17	19	19	19	19
N	707	860	865	868	868

Notes: This table shows the OLS regressions corresponding to the main 2SLS regressions in Table 2.2. Apart from that, the notes to Table 2.2 apply.

B.5.4 IV Regressions, First-Stage and OLS Regressions Using Quota IV

The following tables repeat the estimations from the main part of this study using Bloom et al.'s (2016) main instrument, which is based on the abolition of quotas for textile and apparel after China's WTO accession, instead of initial conditions.

Table B.12: Similarity (Quota IV)

	(1)	(2)	(3)	(4)	(5)
	Δ sim(comp.)	Δ sim(SIC4)	Δ sim(SIC3)	Δ sim(SIC2)	Δ sim(all)
	b/se	b/se	b/se	b/se	b/se
Δ ShareImChina	2.501	2.744	1.449	1.715	2.048
	(3.636)	(1.884)	(2.393)	(2.063)	(1.939)
Year FE	yes	yes	yes	yes	yes
Companies	65	86	87	90	90
SIC4	25	34	35	38	38
SIC3	19	23	23	26	26
SIC2	12	14	14	14	14
N	92	114	115	118	118

Notes: Like Table 2.2, this tables shows estimations of Equation 2.10 using 2SLS, but instruments for Chinese import competitions with Bloom et al.'s (2016) quota instrument. These quotas affected only textile and apparel, such that the sample is smaller. The sample includes the years 2004 and 2005. Apart from these changes, the notes to Table 2.2 apply.

Table B.13: First Stage of Similarity Regressions with Quota IV

	(1)	(2)	(3)	(4)	(5)
	Δ ShareImChina	Δ ShareImChina	Δ ShareImChina	Δ ShareImChina	Δ ShareImChina
	b/se	b/se	b/se	b/se	b/se
Δ Quota	0.047	0.075	0.117	0.112*	0.112*
	(0.046)	(0.057)	(0.069)	(0.056)	(0.056)
Year FE	yes	yes	yes	yes	yes
Obs.	92	114	115	118	118
R^2	0.05	0.04	0.11	0.12	0.12
F	1.38	0.99	1.43	1.99	1.99

Notes: This table shows the first-stage regressions belonging to the 2SLS regressions in Table B.12, which use the quota instrument.

Table B.14: Similarity (OLS for Quota IV Sample)

	(1)	(2)	(3)	(4)	(5)
	Δ sim(comp.)	Δ sim(SIC4)	Δ sim(SIC3)	Δ sim(SIC2)	Δ sim(all)
	b/se	b/se	b/se	b/se	b/se
Δ ShareImChina	-1.608	-0.121	-0.375	-0.383	-0.148
	(1.598)	(0.451)	(0.514)	(0.491)	(0.618)
Year FE	yes	yes	yes	yes	yes
Companies	65	86	87	90	90
SIC4	25	34	35	38	38
SIC3	19	23	23	26	26
SIC2	12	14	14	14	14
N	92	114	115	118	118

Notes: This table shows the OLS regressions based on the sample in the 2SLS regressions using the quota instrument in Table B.12.

B.5.5 Details and Variations of Latent Semantic Analysis

The main part of the chapter uses a low-rank approximation with 100 components in the LSA. This is based on a commonly applied rule of thumb. The number is found in Deerwester et al. (1990) and recommended in the documentation of the scikit-learn package, for example. The truncated SVD matrix explains about 20% of the variance (see Figure B.7). To demonstrate that the decision to follow this rule is not consequential for the qualitative results of this study, this section lists the same regressions with different variations of the text-based measures.

Table B.15 shows the results based on a truncated SVD matrix with 1000 components for slightly more than 50% explained variation.

Figure B.7: Graph of Explained Variance

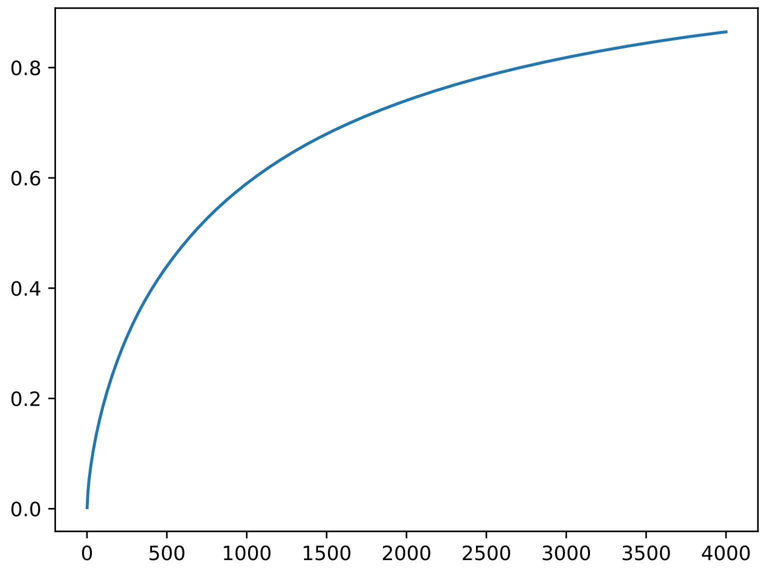

Notes: Graph of explained variance (vertical axis) for singular value decomposition with different numbers of components (horizontal axis)

Table B.15: Similarity (LSA with 1000 Components)

	(1)	(2)	(3)	(4)	(5)
	Δ sim(comp.)	Δ sim(SIC4)	Δ sim(SIC3)	Δ sim(SIC2)	Δ sim(all)
	b/se	b/se	b/se	b/se	b/se
Δ ShareImChina	1.399***	0.193	0.339	0.246	0.192
	(0.528)	(0.263)	(0.237)	(0.188)	(0.196)
Year FE	yes	yes	yes	yes	yes
Companies	298	421	425	428	428
SIC4	114	135	139	142	142
SIC3	67	74	75	78	78
SIC2	17	19	19	19	19
N	707	860	865	868	868

Notes: This table estimates the regressions from Table 2.2 with a similarity measure that is based on LSA with 1000 components (instead of 100). The notes to Table 2.2 apply.

SVD for dimensionality reduction is known to be able to improve the identification of concepts from texts compared to the original tf-idf matrix. The qualitative results are robust to skipping this step and directly comparing tf-idf vectors. Table B.16 shows the results with the 1000 most common terms. (As above, this happens after terms that appear in more than half of all documents are removed.)

Table B.16: Similarity (tf-idf Only)

	(1)	(2)	(3)	(4)	(5)
	Δ sim(comp.)	Δ sim(SIC4)	Δ sim(SIC3)	Δ sim(SIC2)	Δ sim(all)
	b/se	b/se	b/se	b/se	b/se
Δ ShareImChina	1.248**	0.161	0.301	0.205	0.203
	(0.507)	(0.254)	(0.244)	(0.199)	(0.214)
Year FE	yes	yes	yes	yes	yes
Companies	298	421	425	428	428
SIC4	114	135	139	142	142
SIC3	67	74	75	78	78
SIC2	17	19	19	19	19
N	707	860	865	868	868

Notes: This table estimates the regressions from Table 2.2 with a similarity measure based on the cosine similarity of tf-idf vectors with 1000 most common terms (without singular value decomposition). The notes to Table 2.2 apply.

B.5.6 Topic Modeling with Latent Dirichlet Allocation

A generative probabilistic model that can be applied to text corpora and that is increasingly used in the literature is the latent Dirichlet allocation (Blei et al., 2003). I follow Kaplan and Vakili (2015) in estimating the model with 100 topics. The results are shown in Table B.17. Since a probability distribution over topics is estimated for each document, in addition to cosine similarity, the table's second set of columns uses the (five-year difference in) Kullback-Leibler divergence. A decrease in this divergence implies more similar patents, i.e., the results point in the same direction and are in line with the ones reported in the main part of the chapter. The hyperparameters of the Dirichlet priors are both set to $\frac{1}{100}$ (i.e., $\frac{1}{\text{nr. of topics}}$, which is the scikit-learn default), and the maximum number of iterations has been set to 100.

Table B.17: Similarity and Distributional Divergence Using Latent Dirichlet Allocation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Δ sim(comp.)	Δ sim(SIC4)	Δ sim(SIC3)	Δ sim(SIC2)	Δ sim(all)	Δ KLDiv(comp.)	Δ KLDiv(SIC4)	Δ KLDiv(SIC3)	Δ KLDiv(SIC2)	Δ KLDiv(all)
	b/se	b/se	b/se	b/se	b/se	b/se	b/se	b/se	b/se	b/se
Δ ShareImChina	1.609**	0.191	0.196	0.049	0.048	-8.378**	-2.677	-2.940	-1.943	-1.339
	(0.686)	(0.303)	(0.268)	(0.221)	(0.205)	(3.682)	(1.882)	(1.842)	(1.522)	(1.539)
Year FE	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Companies	298	421	425	428	428	298	421	425	428	428
SIC4	114	135	139	142	142	114	135	139	142	142
SIC3	67	74	75	78	78	67	74	75	78	78
SIC2	17	19	19	19	19	17	19	19	19	19
N	707	860	865	868	868	707	860	865	868	868

Notes: The estimates in this table replicate the regressions from Table 2.2 with a similarity measure based on LDA with 100 topics. A positive similarity corresponds to a negative Kullback-Leibler divergence, which is a more suitable comparison for the probability distributions estimated in this statistical model. The notes to Table 2.2 apply.

Appendix C

Appendix to Chapter 3

C.1 Appendix to Section 3.2

C.1.1 Compulsorily Licensed Patents by NBER Technological Subcategory

Figure C.1: Compulsorily Licensed Patents by NBER Technological Subcategory

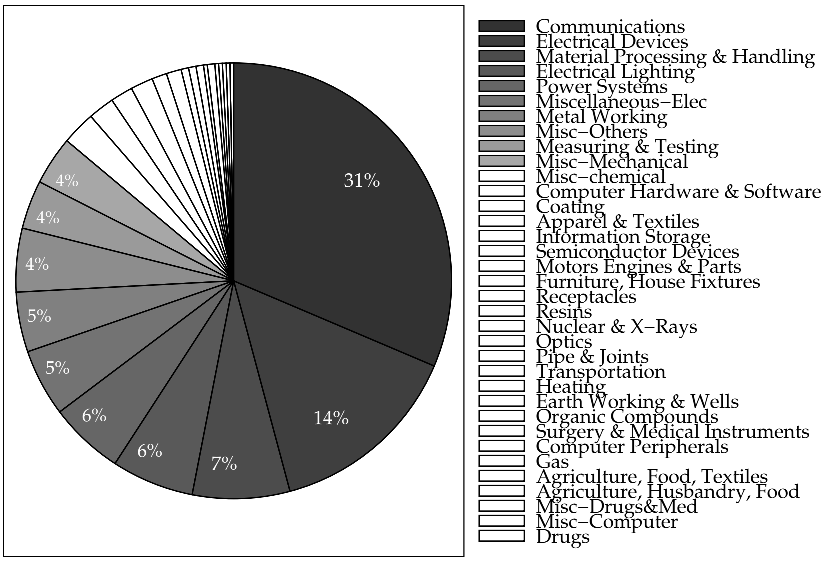

Notes: The pie chart shows the distribution of compulsorily licensed patents over 35 NBER technological subcategories. The legend is sorted from largest share to smallest. The categorization in technological subcategories is based on US patent classifications, following Hall et al. (2001). The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

C.1.2 Patenting of Bell in Radar and Cryptography

Figure C.2: War Technologies Created by Bell Labs

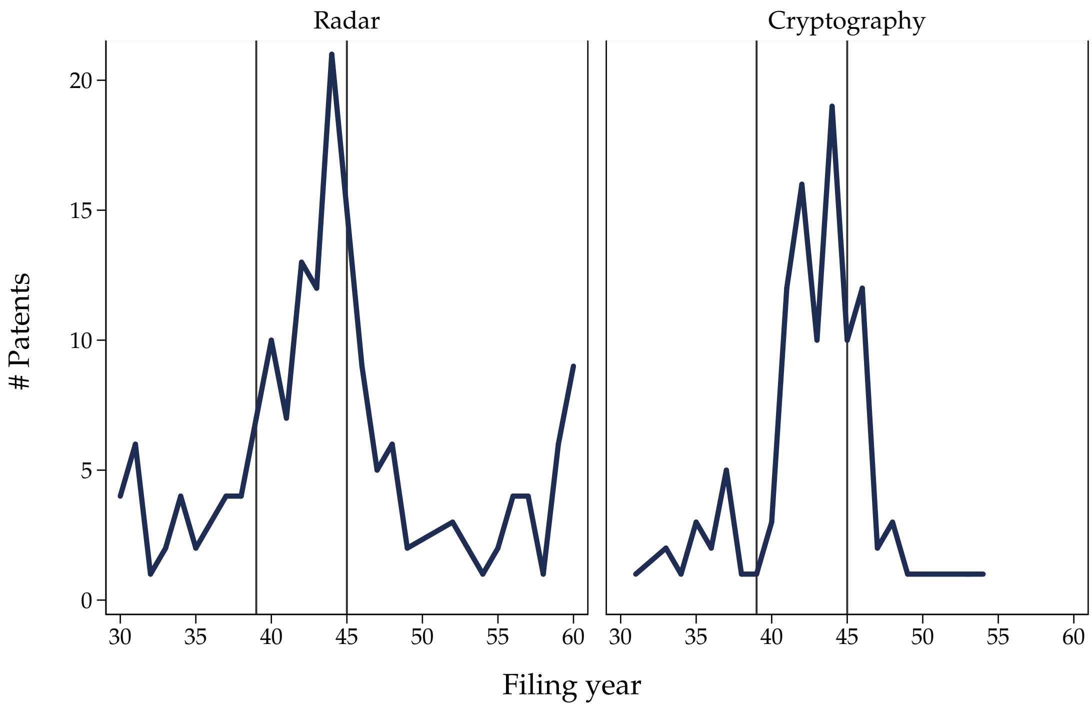

Notes: This figure shows the yearly number of Bell patents relating to radar and cryptography, two technologies relevant for World War II. We identify both technologies by their USPC class: We use the class 342 titled "Communications: directive radio wave systems and devices (e.g., radar, radio navigation)" to classify radar and class 380 titled "Cryptography" to classify cryptography. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

C.2 Appendix to Section 3.4

C.2.1 Comparing the Average Number of Citations of Treatment and Control Patents

In Figure C.3 we compare the evolution of patent citations to Bell patents and control patents in the same publication year and the same four-digit technology class. We use the weights proposed by Iacus et al. (2009) to adjust for the different number of control patents for each Bell patent. From 1949 to 1953, the average number of citations of treatment and control

Figure C.3: Average Number of Citations to Bell and Control Patents Published before 1949

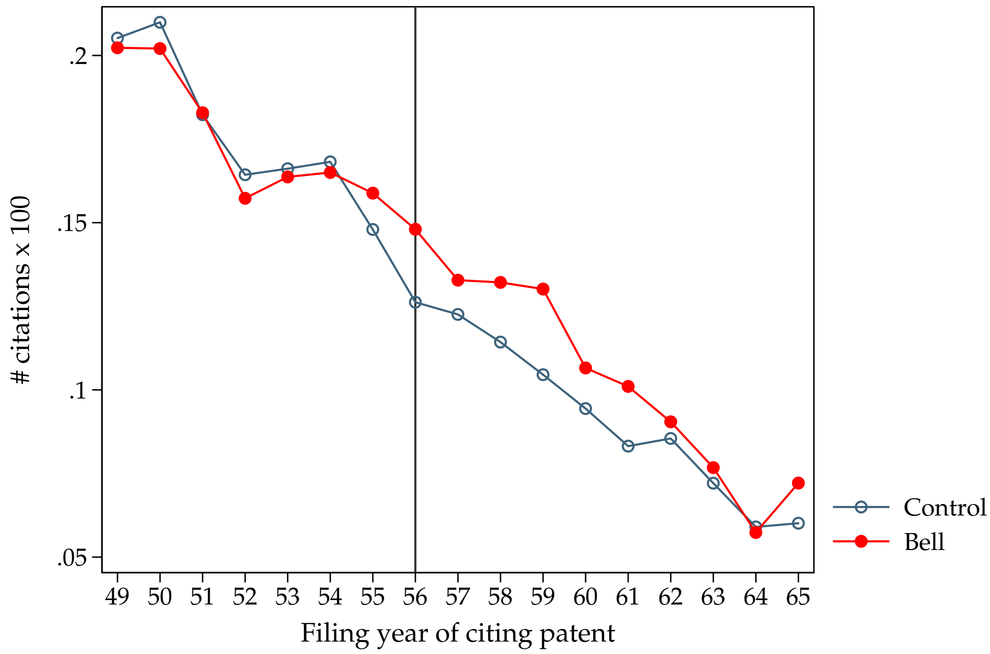

Notes: This figure shows average patent citations of patents published before 1949 in every year after publication. The line with solid circles shows patent citations of the treated patents (Bell patents) and the line with empty circles shows patent citations of control patents, with the same publication year and the same four-digit technology class as the Bell patents. For aggregation we use the weights of Iacus et al. (2009) to adjust for a different number of control patents for each Bell patent. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

patents track each other very closely. This implies that the Bell patents and the control patents exhibit a parallel trend in citations in the first four years after the plea. The two lines diverge in 1954, with Bell patents receiving relatively more citations than control patents, and they converge again in 1961/1962. This is prima facie evidence for an effect from 1954 onward.

C.2.2 Pseudo Outcomes: Unaffected Companies Have No Excess Citations

In the main part of the text we use time varying coefficients to show that there are no yearly excess citations from the B-2 companies, which were exempt from the compulsory licensing agreement. In column 2 of Table C.1 we estimate the average effect for these companies and find none. There are also two other groups of companies that were to a lesser degree affected by the consent decree: foreign companies and companies that already had licensing agreements in place.[1] Foreign companies could license for free but did not receive any technical description or assistance from Bell.[2] In Table C.1 we show the results using as the dependent variable the citations from foreign companies in column 3 and from companies that had a license before the consent decree in column 4. In the last column we use data on all companies that did not have a license from Bell. We do not find a measurable effect for foreign companies or companies with a license and a large effect for companies without a license.

[1] All companies with a license agreement are listed in the hearing documents (Antitrust Subcommittee, 1958, p. 2758).

[2] Verbatim in the consent decree "The defendants are each ordered and directed (...) to furnish to any person domiciled in the United States and not controlled by foreign interests (...) technical information relating to equipment (...)".

Table C.1: The Effect of Compulsory Licensing on Subsequent Citations of Unaffected Companies

	(1) Baseline	(2) B-2 Companies	(3) Foreign companies	(4) License	(5) No license
Treatment	-0.4	-0.1	-0.0	0.5***	-0.9**
	(0.5)	(0.2)	(0.1)	(0.2)	(0.4)
I(55-60)	-6.4***	-1.2***	2.1***	-1.1***	-5.4***
	(0.6)	(0.2)	(0.3)	(0.2)	(0.5)
T x I(55-60)	2.0***	0.2	-0.0	0.4	1.6***
	(0.6)	(0.1)	(0.2)	(0.3)	(0.5)
Constant	18.3***	2.3***	0.9***	3.1***	15.2***
	(1.2)	(0.3)	(0.1)	(0.3)	(1.0)
# treated	4533	4598	4533	4533	4533
Clusters	225	225	225	225	225
Obs.	896556	1096212	896556	896556	896556

Notes: This table shows the results from a difference-in-differences estimation with years 1949-1954 as pre-treatment period and 1955-1960 as treatment period. The estimation equation is

$$\#Citations_{i,t} = \beta_1 \cdot Bell_i + \beta_2 \cdot I[1955-1960] + \beta_3 \cdot Bell_i \cdot I[1955-1960] + \varepsilon_{i,t} \tag{C.1}$$

where $I[1955-1960]$ is an indicator variable for the treatment period 1955-1960. The variable "Bell" is an indicator variable equal to one if a patent is published by a Bell System company before 1949 and therefore treated by the consent decree. As dependent variable we use in the first column all citations by companies other than the filing company. In the second column we use all citations of companies exempt from the consent decree (GE, RCA, Westinghouse & ITT) and in the third column all citations of foreign companies. In the fourth column we use citations of companies that had no licensing agreement with any Bell company prior to the consent decree and in the last column we look at the citation of companies that had a licensing agreement. As control patents, we use all patents that were published in the U.S. matched by publication year, primary United States Patent Classification (USPC) technology class and the number of citations up to 1949. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office. All coefficients are multiplied by 100 for better readability. Standard errors are clustered on the three-digit USPC technology class level and *, **, *** denote statistical significance on 10%, 5% and 1% level, respectively.

C.2.3 Pseudo Treatment: Citation Substitution is Small

One possible interpretation of our estimates is that due to the free availability of Bell technology, companies substituted away from other, potentially more expensive technologies. If this were the case, we should find a negative impact of the consent decree on citations of similar patents of other companies.[3] To see if this is the case, we assign a pseudo treatment to the patents of GE, RCA, Westinghouse, which were part of the B-2 agreement, and ITT. These companies were among the largest patenting firms in the ten technology classes in which Bell had most patents between 1939 and 1949. Results are reported in Table C.2, column 2. We find no effect, implying that the citation substitution is either small or homogeneous to patents of these companies and the control group.

For a second approach, we exploit the fact that a patent's technology is classified twice: once in the USPC system, which has a technical focus, and once in the IPC system, which reflects more closely the intended industry or profession ("usage") (Lerner, 1994). In columns 3 and 4 of Table C.2 we assign a pseudo-treatment to all patents that have the same USPC class and the same IPC class as the Bell patents. As control group we use in column 3 patents with the same USPC, but a different IPC classification as Bell patents. In column 4 we use as a control group patents with the same IPC, but a different USPC classification as Bell patents. Thus we compare patents that are arguably more similar to the Bell patents to two different control groups. We find a small, negative but statistically insignificant effect. Again, this speaks in favor of limited citation substitution or - alternatively - a homogeneous citation substitution to all control groups.

[3] This approach is suggested by Imbens and Rubin (2015).

Table C.2: Auxiliary Regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Pseudo Treatment			Diff. Control Group		
	Base-line	B-2 Comp.	Control: Same USPC, diff IPC	Control: Same IPC, diff USPC	Control: Same IPC, diff USPC	Control: Same IPC	Loose
Treatment	-0.4	-0.6	0.7**	-0.1	-0.3	0.4	0.7
	(0.5)	(0.4)	(0.3)	(0.5)	(1.6)	(0.7)	(0.8)
I(55-60)	-6.4***	-4.2***	-1.5***	-2.2***	-6.4***	-6.0***	-6.7***
	(0.6)	(0.4)	(0.5)	(0.2)	(0.7)	(0.5)	(0.6)
T x I(55-60)	2.0***	0.0	-0.5	-0.3	1.2	1.5**	1.5**
	(0.6)	(0.4)	(0.5)	(0.4)	(0.8)	(0.6)	(0.6)
Constant	18.3***	16.6***	11.9***	13.3***	19.7***	17.8***	18.6***
	(1.2)	(0.6)	(0.5)	(0.3)	(1.0)	(1.0)	(1.0)
# treated	4533	7869	42607	48526	4649	4511	4665
Clusters	225	207	202	398	398	386	230
Obs.	896556	836172	760452	1348440	835188	705612	1301412

Notes: This table shows the results from a difference-in-differences estimation with years 1949-1954 as pre-treatment period and 1955-1960 as treatment period, controlling for year fixed effects. As dependent variable, we use all citations by companies other than the filing company. As control patents, we use all patents that were published in the U.S. matched by publication year, primary USPC technology class, and the number of citations up to 1949. In all columns we match the control patents on publication year and the number of citations prior to 1949. In columns 2 to 4 we assign pseudo treatments. In column 2 we define patents of the B-2 companies (GE, RCA, Westinghouse & ITT) as treated and match the control patents on the USPC class. In column 3 we assign all patents that have the same USPC and different three-digit IPC technology class as treated and in column 4 we assign patents with the same IPC and different USPC classification as treated. In column 5 we use as controls patents in the same three-digit IPC class but in a different USPC class than the Bell patents. In column 6 we use as controls patents with the same four-digit IPC class as the Bell patents. In column 7 we coarsen the publication year to two year windows and sort all pre-citations in 10 equally sized bins to match a larger number of patents. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office. All coefficients are multiplied by 100 for better readability. Standard errors are clustered on the three-digit USPC technology class level and *, **, *** denote statistical significance on 10%, 5% and 1% level, respectively.

C.2.4 Effects are Robust to Different Matching Strategies

In columns 5 to 7 of Table C.2 and in Figure C.4 we report results from using several alternative matching variables. In the main specification, we use the age (measured by the publication year), the technology (measured by USPC class) and the quality of a patent (measured by the number of citations up to 1949). In column 6 we use patents in the same IPC but different USPC class instead of using those in the same USPC class. In column 7 we match on the IPC classification, independent of the USPC class. Finally, in column 8 we do a coarsened exact matching in order to match all Bell patents.[4] In all three cases the size of the effects is similar to the one in the main specification. In Figure C.4 we show the size of the treatment effects for different combinations of background variables as proxy for age, technology and quality. On the vertical axis we plot the number of matched patents. The coefficient is mostly around 2.

[4] Coarsened exact matching was proposed by Iacus et al. (2012). In this specification we match on one of five publication year categories that contain two years each and one of ten prior-citation categories.

Figure C.4: Treatment Effects for Different Matching Variables

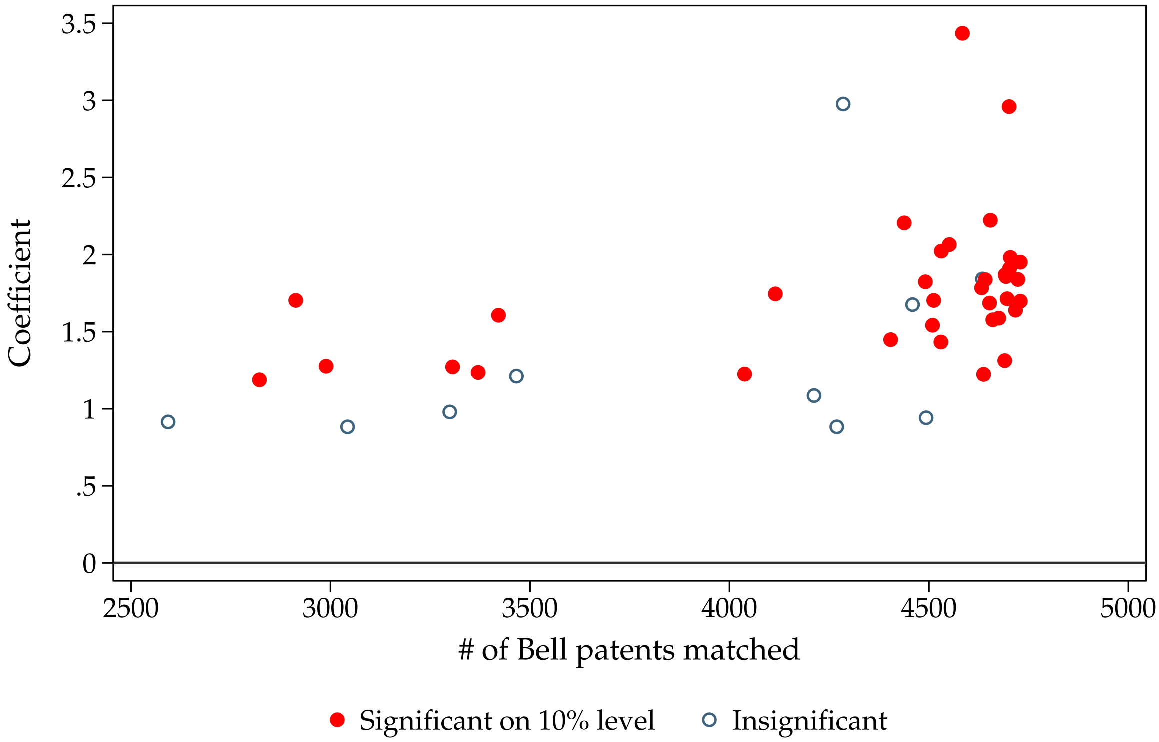

Notes: In this figure we plot the parameter estimates from difference-in-differences estimations of the impact of the consent decree for different matching strategies, controlling for year fixed effects. As before, as dependent variable we use all citations by companies other than the filing company. In all regressions, we use a measure for the age, the technology and the quality of a patent for matching. As measures for the age of a patent, we alternatively use application year, publication year or both. For technology, we use the USPC, the USPC with subclasses, the three and the four-digit IPC. As a measure of quality, we use the number of pre-citations as exact numbers, coarsened to steps of five citations and an indicator for at least one citation prior to 1949. The horizontal axis displays the number of matched Bell patents. Empty symbols are insignificant and full symbols are significant at the 10% level. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

C.2.5 Patenting Behavior of Bell Relative to Comparable Companies

Figure C.5: Patenting of Bell System and B-2 Companies without RCA

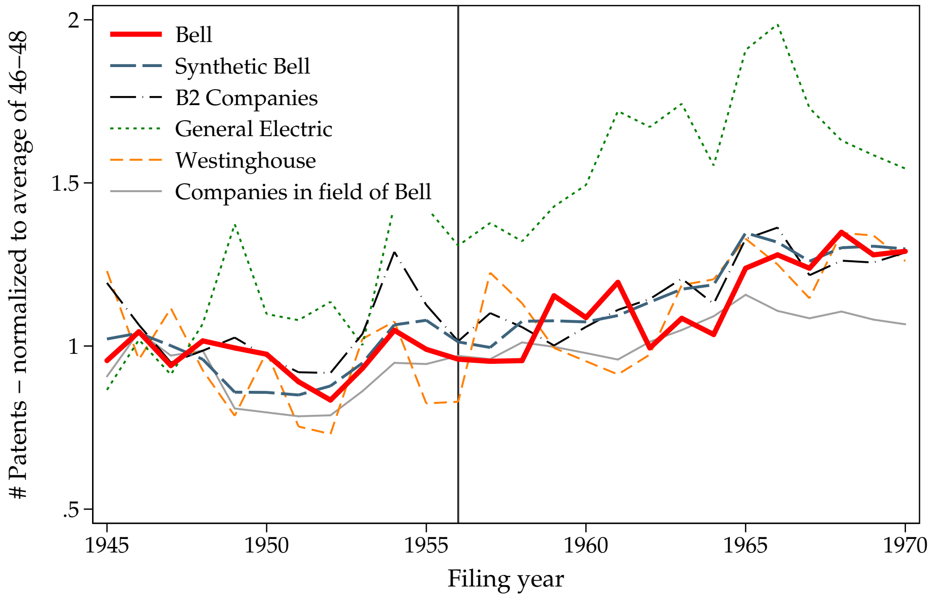

Notes: In this figure we compare Bell's total patenting to a synthetic Bell, the number of patents filed by the B-2 companies (General Electric, Westinghouse, RCA and ITT), General Electric and Westinghouse separately and all companies that existed before 1949 and had at least 100 patents in any field in which Bell was active. The number of patents are normalized to the average number of patents from 1946-1948. We show General Electric and Westinghouse separately, because RCA had a consent decree involving patents in 1958 and thus might have changed its behavior. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

C.2.6 Share of Communication Patents Measured with NBER Technology Subcategories

Figure C.6: Share of Communication Patents

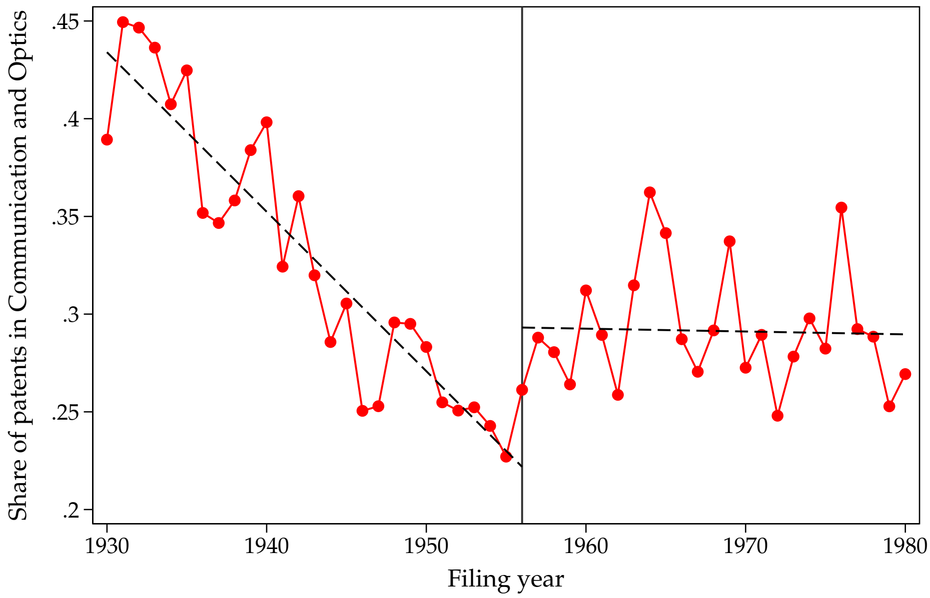

Notes: This figure shows the share of patents related to communication relative to all patents filed by Bell. We define technologies related to communication as the NBER subcategories "Communication" and "Optics" (Hall et al., 2001). We include "Optics" because after the invention of the laser at Bell Labs in 1958, Bell officials predicted correctly that optics might be crucial for the future of communication (Gertner, 2012, p. 253).

C.2.7 Effect for Different Definitions of Small and Young Assignees

In Figure C.7 we estimate the main treatment coefficient separately for citations of different size and age groups of assignees. We find that the effect is driven mainly by companies and individual inventors without patents before 1949 and companies and individual inventors that have been active for less than one year at the time of the citations.

Figure C.7: Sample Split by Characteristics of Citing Firm

(a) By size of patent portfolio in 1949

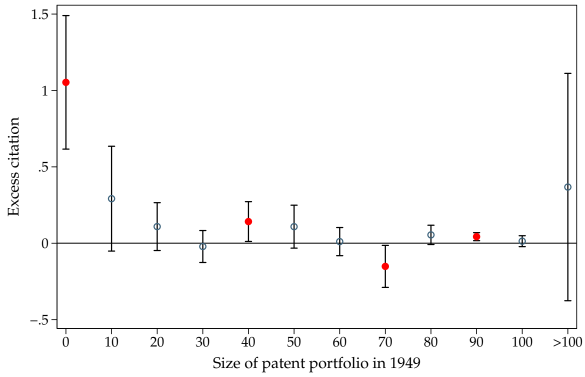

(b) By age of company at citation

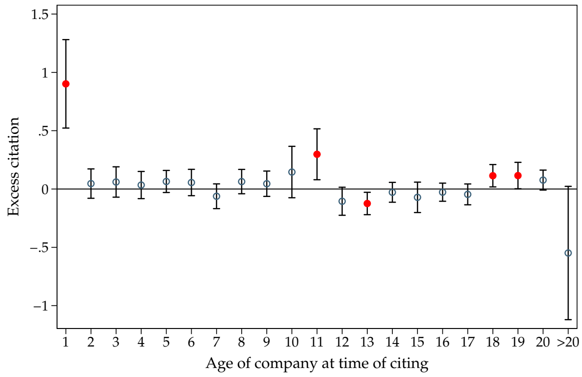

Notes: These subfigures show results from a difference-in-differences estimation with the years 1949-1954 as pre-treatment period and 1955-1960 as treatment period, controlling for year fixed effects. As dependent variable we use all citations by companies other than the filing companies with a specific size of their patent portfolio (Subfigure (a)) and a specific company age (b) as indicated in the figure. As control patents we use all patents that were published in the U.S. matched by publication year, primary USPC technology class, and the number of citations up to 1949. The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

C.3 Appendix to Section 3.5

C.3.1 Effect by NBER Technology Subcategory

In this section we estimate our main treatment effect separately for citations of patents in different NBER technology subcategories. The results are reported in Figure C.8. The increase in citations comes mainly from technologies related to electrical components, in particular in "Electrical Devices". Yet, there is no increase in citations by patents in the subcategory of "Communication". These results corroborate the finding in our main text that there is no increase in follow-on innovation in industries concerned with production of communication equipment, the core business of Bell.

Figure C.8: Effect of Compulsory Licensing on Subsequent Citations By NBER Technological Subcategory

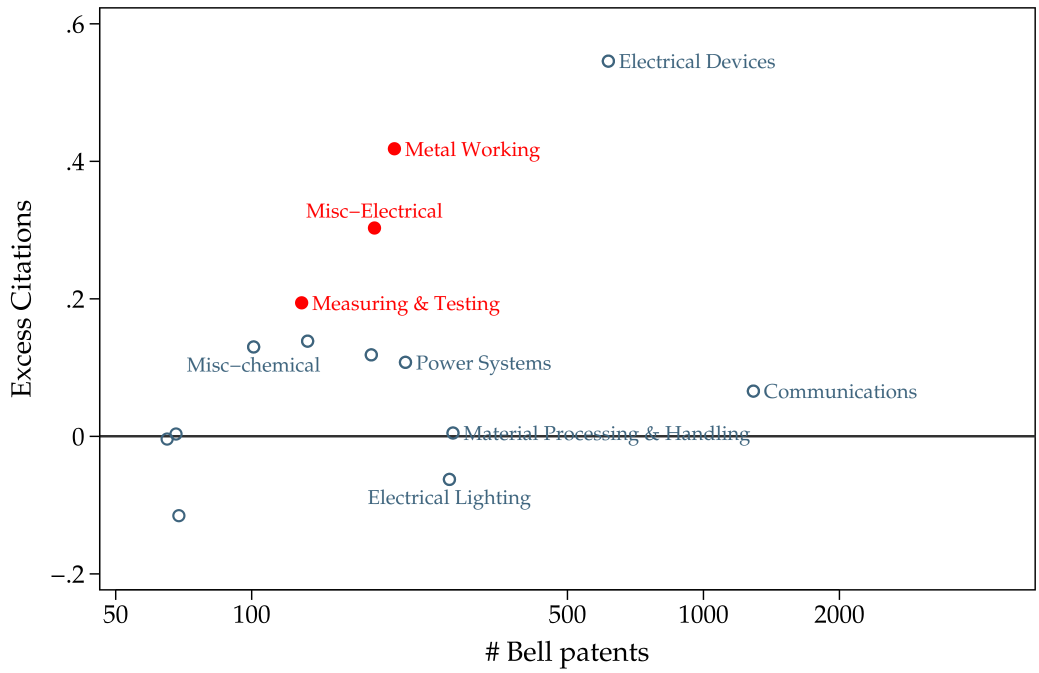

Notes: This figure shows difference-in-differences estimates of the impact of the consent decree on citations from patents in different NBER technological subcategories, controlling for year fixed effects. As dependent variable we use all citations by companies other than the filing company. As control patents we use all patents that were published in the U.S., matched by publication year, primary USPC technology class, and the number of citations up to 1949. A solid circle means that the coefficient is significant at the 10% level. We split the citing patents by NBER technology subcategory following Hall et al. (2001). The data are from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

C.3.2 No Lack of Follow-On Innovation in Telecommunications

This section presents evidence that the null effect in telecommunications was not due to a lack in potential follow-on innovation in the telecommunications market. To do this we look at the total number of citations, the sum of citations of other companies and self-citations, to Bell patents inside and outside of telecommunications. In Subfigure (a) of Figure C.9 we plot the average number of total citations to Bell patents related to communication and related to other fields. We use the concordance of Kerr (2008) to assign to each Bell patent the most likely SIC code. We find that the total number of citations to telecommunications patents of Bell was at least as high as to patents outside of communication. This speaks against a low quality of compulsorily licensed patents as a reason for the lack in follow-on innovation in telecommunications. In Subfigure (b) we show that the total number of patent citations to Bell's patents inside and outside of telecommunications were also almost identical before and after the consent decree. This suggests that after the consent decree the potential for follow-on innovation was not significantly lower in telecommunications than in other fields.

Figure C.9: Number of Citations to Bell Patents Inside and Outside of Communication

(a) Up to 1949 (b) All Bell Patents

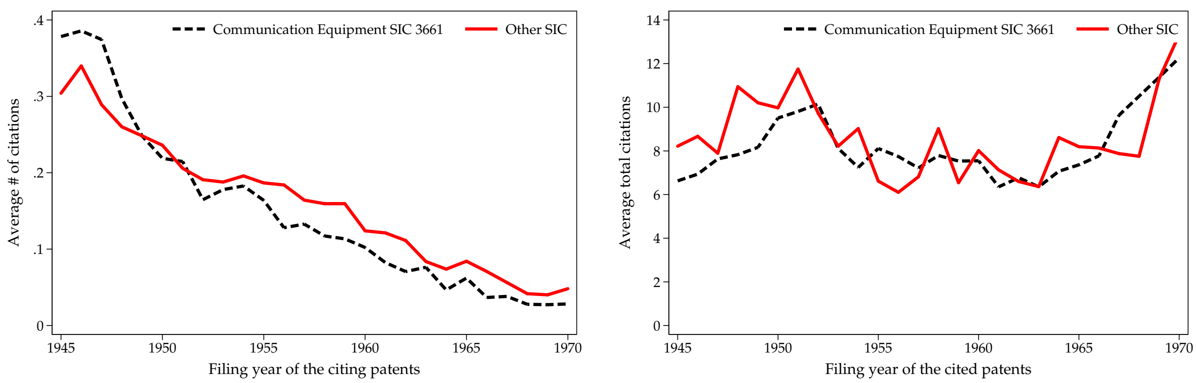

Notes: Subfigure (a) shows the average number of citations per year for all Bell patents that are most likely used in the production of communication equipment (SIC 3661) and that are used in any other industry. To classify a patent by its most likely industry, we use the data of Kerr (2008) to assign to each USPC class the most likely four-digit SIC industry in which it is used. Subfigure (b) shows the total number of citations to Bell patents inside and outside of telecommunication filed in a particular year. In this graph we use total citations, the sum of citations from other companies and from Bell to its own patents. The data stem from the Worldwide Patent Statistical Database (PATSTAT) of the European Patent Office.

List of Figures

List of Tables

Bibliography

Acemoglu, D. and Akcigit, U. (2012). Intellectual property rights policy, competition and innovation. *Journal of the European Economic Association*, 10(1):1–42.

Acemoglu, D., Autor, D., Dorn, D., Hanson, G. H., and Price, B. (2015). Import competition and the great US employment sag of the 2000s. *Journal of Labor Economics*, 34(S1):S141–S198.

Aghion, P., Bloom, N., Blundell, R., Griffith, R., and Howitt, P. (2002). Competition and innovation: an inverted U relationship. *IFS Working Paper (W02/04)*.

Aghion, P., Bloom, N., Blundell, R., Griffith, R., and Howitt, P. (2005). Competition and innovation: An inverted-U relationship. *The Quarterly Journal of Economics*, 120(2):701–728.

Aghion, P., Harris, C., Howitt, P., and Vickers, J. (2001). Competition, imitation and growth with step-by-step innovation. *The Review of Economic Studies*, 68(3):467–492.

Aghion, P. and Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, 60(2):323–351.

Aghion, P. and Howitt, P. W. (2008). *The economics of growth*. MIT press.

Akcigit, U. and Kerr, W. R. (2010). Growth through heterogeneous innovations. *NBER Working Paper No. 16443*.

Alcacer, J. and Gittelman, M. (2006). Patent citations as a measure of knowledge flows: The influence of examiner citations. *The Review of Economics and Statistics*, 88(4):774–779.

Alcacer, J., Gittelman, M., and Sampat, B. (2009). Applicant and examiner citations in US patents: An overview and analysis. *Research Policy*, 38(2):415–427.

Amiti, M. and Freund, C. (2010). The anatomy of china's export growth. In *China's growing role in world trade*, pages 35–56. University of Chicago Press.

Angrist, J. D. and Pischke, J.-S. (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.

Antitrust Subcommittee (1958). Consent decree program of the Department of Justice. American Telephone & Tel. Subcommittee - hearings before the Antitrust Subcommittee. *U.S. House Committee on the Judiciary, Committee on the Judiciary Serial No. 9*.

Antitrust Subcommittee (1959). Consent decree program of the Department of Justice. *Committee on the Judiciary, House of Representatigves, Eighty-Sixth Congress, First Session*.

Arora, A., Belenzon, S., and Patacconi, A. (2015). Killing the golden goose? The decline of science in corporate R&D. *NBER Working Paper No. 20902*.

Arrow, K. (1962). Economic welfare and the allocation of resources for invention. In *The rate and direction of inventive activity: Economic and social factors*, pages 609–626. Princeton University Press.

Autor, D., Dorn, D., and Hanson, G. H. (2013). The China syndrome: Local labor market effects of import competition in the United States. *The American Economic Review*, 103(6):2121–2168.

Autor, D., Dorn, D., Hanson, G. H., Pisano, G., and Shu, P. (2016a). Foreign competition and domestic innovation: Evidence from U.S. patents. *NBER Working Paper No. 22879*.

Autor, D. H., Dorn, D., and Hanson, G. H. (2016b). The China Shock: Learning from labor-market adjustment to large changes in trade. *Annual Review of Economics*, 8:205–240.

Baker, J. B. (2012). Exclusion as a core competition concern. *Antitrust Law Journal*, 78:527–589.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636.

Balsmeier, B., Fleming, L., and Manso, G. (2017). Independent boards and innovation. *Journal of Financial Economics*, 123(3):536–557.

Balsmeier, B., Li, G.-C., Chesebro, T., Zang, G., Fierro, G., Johnson, K., Kaulagi, A., Lück, S., O'Reagan, D., Yeh, B., and Others (2016). Machine learning and natural language processing on the patent corpus: data, tools, and new measures. *Working Paper*.

Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)*, pages 289–300.

Bholat, D., Hans, S., Santos, P., and Schonhardt-Bailey, C. (2015). *Text mining for central banks*. Number 33 in Handbooks. Centre for Central Banking Studies, Bank of England.

Bird, S., Klein, E., and Loper, E. (2009). *Natural language processing with Python: analyzing text with the natural language toolkit*. O'Reilly Media, Inc.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan):993–1022.

Bloom, N., Draca, M., and Van Reenen, J. (2016). Trade induced technical change? The impact of chinese imports on innovation, IT and productivity. *Review of Economic Studies*, 83(1):87–117.

Bosetti, V., Cattaneo, C., and Verdolini, E. (2015). Migration of skilled workers and innovation: A European perspective. *Journal of International Economics*.

Braunerhjelm, P., Ding, D., and Thulin, P. (2015). Does labour mobility foster innovation? Evidence from Sweden. Working Paper Series in Economics and Institutions of Innovation 403, Royal Institute of Technology, CESIS - Centre of Excellence for Science and Innovation Studies.

Cabral, L. (1994). Bias in market R&D portfolios. *International Journal of Industrial Organization*, 12(4):533–547.

Cabral, L. M. B. (2003). R&D competition when firms choose variance. *Journal of Economics Management Strategy*, 12(1):139–150.

Cameron, A. C. and Trivedi, P. K. (2005). *Microeconometrics: methods and applications*. Cambridge University Press.

Choudhury, P. (2015). Return migration and geography of innovation in MNEs: A natural experiment of on-the-job learning of knowledge production by local workers reporting to return migrants. *Journal of Economic Geography*, 16(3):585–610.

Dauth, W., Findeisen, S., and Suedekum, J. (2014). The rise of the east and the far east: German labor markets and trade integration. *Journal of the European Economic Association*, 12(6):1643–1675.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American society for information science*, 41(6):391.

Delrahim, M. (2004). Forcing firms to share the sandbox: Compulsory licensing of intellectual property rights and antitrust. *European Business Law Review*, 15(5):1059–1069.

Doran, K., Gelber, A., and Isen, A. (2015). The effect of high-skilled immigration policy on firms: Evidence from H-1B visa lotteries. *NBER Working Paper No. 20668*.

Duguet, E. and MacGarvie, M. (2005). How well do patent citations measure flows of technology? Evidence from French innovation surveys. *Economics of Innovation and New Technology*, 14(5):375–393.

Eberhardt, M., Helmers, C., and Yu, Z. (2016). What can explain the Chinese patent explosion? *Oxford Economic Papers*, 69(1):239–262.

Ederer, F. (2013). Incentives for parallel innovation. *Working Paper*.

Ederer, F. and Manso, G. (2013). Is pay for performance detrimental to innovation? *Management Science*, 59(7):1496–1513.

Foster, L., Haltiwanger, J., and Krizan, C. J. (2006). Market selection, reallocation, and restructuring in the US retail trade sector in the 1990s. *The Review of Economics and Statistics*, 88(4):748–758.

Galasso, A. (2012). Broad cross-license negotiations. *Journal of Economics & Management Strategy*, 21(4):873–911.

Galasso, A. and Schankerman, M. (2015a). Patent rights and innovation by small and large firms. *NBER Working Paper No. 21769*.

Galasso, A. and Schankerman, M. (2015b). Patents and cumulative innovation: Causal evidence from the courts. *The Quarterly Journal of Economics*, 130(1):317–369.

Gans, J. S., Hsu, D. H., and Stern, S. (2002). When does start-up innovation spur the gale of creative destruction? *The RAND Journal of Economics*, 33(4):571–586.

Gans, J. S., Hsu, D. H., and Stern, S. (2008). The impact of uncertain intellectual property rights on the market for ideas: Evidence from patent grant delays. *Management Science*, 54(5):982–997.

Gans, J. S. and Stern, S. (2003). The product market and the market for ideas: commercialization strategies for technology entrepreneurs. *Research policy*, 32(2):333–350.

Gentzkow, M., Kelly, B. T., and Taddy, M. (2017). Text as data. *NBER Working Paper No. 23276*.

Gertner, J. (2012). *The idea factory: Bell Labs and the great age of American innovation*. Penguin.

Gilbert, R. (2006). Looking for Mr. Schumpeter: Where are we in the competition-innovation debate? *Innovation policy and the economy*, 6:159–215.

Green, J. R. and Scotchmer, S. (1995). On the division of profit in sequential innovation. *The RAND Journal of Economics*, 26(1):20–33.

Griliches, Z. (1979). Issues in assessing the contribution of research and development to productivity growth. *Bell Journal of Economics*, 10(1):92–116.

Grindley, P. C. and Teece, D. J. (1997). Licensing and cross-licensing in semiconductors and electronics. *California Management Review*, 39(2):8–41.

Grossman, G. M. and Helpman, E. (1994). Endogenous innovation in the theory of growth. *The Journal of Economic Perspectives*, 8(1):23–44.

Hall, B. H., Jaffe, A. B., and Trajtenberg, M. (2001). The NBER patent citation data file: Lessons, insights and methodological tools. *NBER Working Paper No. 8498*.

Hansen, S. and McMahon, M. (2016). Shocking language: Understanding the macroeconomic effects of central bank communication. *Journal of International Economics*, 99:S114—-S133.

Hansen, S., McMahon, M., and Prat, A. (2017). Transparency and deliberation within the FOMC: a computational linguistics approach. *Working Paper*.

Harhoff, D., Narin, F., Scherer, F. M., and Vopel, K. (1999). Citation frequency and the value of patented inventions. *The Review of Economics and Statistics*, 81(3):511–515.

Helpman, E. (1998). *General purpose technologies and economic growth*. MIT press.

Hoberg, G. and Phillips, G. (2010). Product market synergies and competition in mergers and acquisitions: A text-based analysis. *The Review of Financial Studies*, 23(10):3773–3811.

Holbrook, D., Cohen, W. M., Hounshell, D. A., and Klepper, S. (2000). The nature, sources, and consequences of firm differences in the early history of the semiconductor industry. *Strategic Management Journal*, 21(10-11):1017–1041.

Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, pages 65–70.

Hu, A. G. and Jaffe, A. B. (2003). Patent citations and international knowledge flow: the cases of Korea and Taiwan. *International Journal of Industrial Organization*, 21(6):849–880.

Hu, A. G. and Jefferson, G. H. (2008). Science and technology in China. In Brandt, L. and Rawski, T. G., editors, *China's great economic transformation*, chapter 9, pages 286–336. Cambridge University Press.

Hu, A. G. and Jefferson, G. H. (2009). A great wall of patents: What is behind China's recent patent explosion? *Journal of Development Economics*, 90(1):57–68.

Hudson, P. (2014). *The Industrial Revolution*. Bloomsbury Publishing.

Iacus, S. M., King, G., and Porro, G. (2009). CEM: Software for coarsened exact matching. *Journal of Statistical Software*, 30(9):1–27.

Iacus, S. M., King, G., and Porro, G. (2012). Causal inference without balance checking: Coarsened exact matching. *Political Analysis*, 20:1–24.

Imbens, G. W. and Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical science*. Cambridge University Press.

Jaffe, A. B. (1986). Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits, and market value. *American Economic Review*, 76(5):984–1001.

Jaffe, A. B. (1989). Real effects of academic research. *American Economic Review*, 79(5):957–70.

Jaffe, A. B. and Lerner, J. (2011). *Innovation and its discontents: How our broken patent system is endangering innovation and progress, and what to do about it*. Princeton University Press.

Jaffe, A. B. and Trajtenberg, M. (1999). International knowledge flows: evidence from patent citations. *Economics of Innovation and New Technology*, 8(1-2):105–136.

Jaffe, A. B., Trajtenberg, M., and Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *The Quarterly Journal of Economics*, 108(3):577–98.

Kaiser, U., Kongsted, H. C., and Ronde, T. (2015). Does the mobility of R&D labor increase innovation? *Journal of Economic Behavior & Organization*, 110(336):91–105.

Kaplan, S. and Vakili, K. (2015). The double-edged sword of recombination in breakthrough innovation. *Strategic Management Journal*, 36(10):1435–1457.

Kerr, S. P. and Kerr, W. R. (2015). Global collaborative patents. *NBER Working Paper No. 21735*.

Kerr, W. R. (2008). Ethnic scientific communities and international technology diffusion. *The Review of Economics and Statistics*, 90(3):518–537.

Kerr, W. R. and Lincoln, W. F. (2010). The supply side of innovation: H-1B visa reforms and U.S. ethnic invention. *Journal of Labor Economics*, 28(3):473–508.

Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics*, 132(2):665–712.

Kwon, I. (2010). R&D portfolio and market structure. *The Economic Journal*, 120(543):313–323.

Lacetera, N. and Zirulia, L. (2012). Individual preferences, organization, and competition in a model of R&D incentive provision. *Journal of Economic Behavior and Organization*, 84(2):550–570.

Lanjouw, J. O. and Schankerman, M. (2004). Protecting intellectual property rights: Are small firms handicapped? *Journal of Law and Economics*, 47(1):45–74.

Lerner, J. (1994). The importance of patent scope: An empirical analysis. *The RAND Journal of Economics*, 25:319–333.

MacGarvie, M. (2005). The determinants of international knowledge diffusion as measured by patent citations. *Economics Letters*, 87(1):121–126.

MacGarvie, M. (2006). Do firms learn from international trade? *The Review of Economics and Statistics*, 88(1):46–60.

Malthus, T. R. (1798). *An Essay on the Principle of Population*. Project Gutenberg ebook, https://www.gutenberg.org/ebooks/4239.

Mankiw, N. G. and Whinston, M. D. (1986). Free entry and social inefficiency. *The RAND Journal of Economics Rand Journal of Economics*, 1717511947(1):48–58.

Manning, C., Raghavan, P., and Schütze, H. (2008). *Introduction to Information Retrieval*. Cambridge University Press.

Manso, G. (2011). Motivating innovation. *The Journal of Finance*, 66:1823–1860.

Marx, M., Singh, J., and Fleming, L. (2015). Regional disadvantage? Employee non-compete agreements and brain drain. *Research Policy*, 44(2):394–404.

Moser, P. (2013). Patents and innovation: Evidence from economic history. *Journal of Economic Perspectives*, 27(1):23–44.

Moser, P., Ohmstedt, J., and Rhode, P. W. (2015). Patent citations and the size of the inventive step - evidence from hybrid corn. *NBER Working Paper No. 21443*.

Moser, P. and Voena, A. (2012). Compulsory licensing: Evidence from the trading with the enemy act. *American Economic Review*, 102(1):396–427.

Moser, P., Voena, A., and Waldinger, F. (2014). German jewish emigres and US invention. *American Economic Review*, 104(10):3222–3255.

Mowery, D. C. (2011). Federal policy and the development of semiconductors, computer hardware, and computer software: A policy model for climate change R&D? In *Accelerating*

Energy Innovation: Insights from Multiple Sectors, pages 159–188. University of Chicago Press.

Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press.

Murray, F. and Stern, S. (2007). Do formal intellectual property rights hinder the free flow of scientific knowledge?: An empirical test of the anti-commons hypothesis. *Journal of Economic Behavior & Organization*, 63(4):648–687.

Naughton, B. (2007). *The Chinese economy: Transitions and growth*. MIT press.

Nelson, R. (1962). The link between science and invention: The case of the transistor. In *The rate and direction of inventive activity: Economic and social factors*, pages 549–584. Princeton University Press.

Nelson, R. R. (1959). The simple economics of basic scientific research. *Journal of Political Economy*, 67(3):297–306.

Niebuhr, A. (2010). Migration and innovation: Does cultural diversity matter for regional R&D activity? *Papers in Regional Science*, 89(3):563–585.

Noll, A. (1987). Bell system R&D activities. *Telecommunications Policy*, 11(2):161 – 178.

Nordhaus, W. (1969). An economic theory of technological change. *American Economic Review*, 59(2):18–28.

Ozgen, C., Peters, C., Niebuhr, A., Nijkamp, P., and Poot, J. (2014). Does cultural diversity of migrant employees affect innovation? *International Migration Review*, 48:S377–S416.

Parrotta, P., Pozzoli, D., and Pytlikova, M. (2014). The nexus between labor diversity and firm's innovation. *Journal of Population Economics*, 27(2):303–364.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Peri, G. (2005). Determinants of knowledge flows and their effect on innovation. *The Review of Economics and Statistics*, 87(2):308–322.

Peruzzi, M., Zachmann, G., and Veugelers, R. (2014). Remerge: Regression-based record linkage. *Bruegel Working Paper.*

Porter, M. F. (1980). An algorithm for suffix stripping. *Program*, 14(3):130–137.

Romano, J. P., Shaikh, A. M., and Wolf, M. (2010). Hypothesis testing in econometrics. *Annu. Rev. Econ.*, 2(1):75–104.

Romer, P. M. (1990). Endogenous technological change. *Journal of political Economy*, 98(5, Part 2):S71–S102.

Rosenberg, N. (1990). Why do firms do basic research (with their own money)? *Research Policy*, 19(2):165–174.

Russell, S. S. (1986). Remittances from international migration: A review in perspective. *World Development*, 14(6):677–696.

Sampat, B. and Williams, H. L. (2015). How do patents affect follow-on innovation? Evidence from the human genome. *NBER Working Paper No. 21666.*

Schmidt, E. and Rosenberg, J. (2014). *How Google Works.* Hachette UK.

Schmidt, K. M. (1997). Managerial incentives and product market competition. *The Review of Economic Studies*, 64(2):191.

Schmutzler, A. (2010). Is competition good for innovation? A simple approach to an unresolved question. *Foundations and Trends in Microeconomics*, 5(6):355–428.

Schumpeter, J. A. (1942). *Socialism, Capitalism and Democracy.* Harper and Brothers.

Scotchmer, S. (1991). Standing on the shoulders of giants: Cumulative research and the patent law. *Journal of Economic Perspectives*, 5(1):29–41.

Segal, I. and Whinston, M. D. (2007). Antitrust in innovative industries. *American Economic Review*, 97(5):1703–1730.

Singh, J. and Marx, M. (2013). Geographic constraints on knowledge spillovers: political borders vs. spatial proximity. *Management Science*, 59(9):2056–2078.

Smith, A. (1776). *The Wealth of Nations.* Project Gutenberg ebook, https://www.gutenberg.org/ebooks/3300.

Staiger, D. and Stock, J. H. (1997). Instrumental variables regression with weak instruments. *Econometrica: Journal of the Econometric Society*, pages 557–586.

Sturiale, J. E. (2011). Compulsory licensing of intellectual property as merger remedy: A decision-theoretic approach. *Louisiana Law Review*, 72:605–646.

Temin, P. and Galambos, L. (1987). *The fall of the Bell system: A study in prices and politics.* Cambridge University Press.

Thompson, P. and Fox-Kean, M. (2005). Patent citations and the geography of knowledge spillovers: A reassessment. *American Economic Review*, 95(1):450–460.

Tirole, J. (1988). *The Theory of Industrial Organization.* MIT press.

Trajtenberg, M. (1990). A penny for your quotes: patent citations and the value of innovations. *The Rand Journal of Economics*, pages 172–187.

Waller, S. W. and Sag, M. (2014). Promoting innovation. *Iowa Law Review*, 100:2223–2247.

Wessner, C. W. et al. (2001). *Capitalizing on new needs and new opportunities: Government-industry partnerships in biotechnology and information technologies.* National Academies Press. Chapter 2 available at http://www.nap.edu/read/10281/chapter/11.

Williams, H. L. (2015). Intellectual property rights and innovation: Evidence from health care markets. In *Innovation Policy and the Economy, Volume 16.* University of Chicago Press.

Wu, T. (2012). Taking innovation seriously: Antitrust enforcement if innovation mattered most. *Antitrust Law Journal*, 78(2):313–328.