Conflicts of Interest

Articles on Common Benefits and Private Profits

Inauguraldissertation zur Erlangung des akademischen Grades eines Doktors der Wirtschaftswissenschaften der Universität Mannheim

Johannes Koenen

vorgelegt im Frühjahrs- und Sommersemester 2010

Abteilungssprecher:	Professor Tom Krebs, Ph.D
Referent:	Professor Konrad Stahl, Ph.D.
Korreferent:	Professor Dr. Martin Peitz
Tag der mündlichen Prüfung:	23. Juni 2010

To my one great *discovery*

Acknowledgements

More people have supported me in the course of the past years than could be included in this whole volume, therefore I can only thank a selection of them here.

I was tremendously fortunate in the academic environment I was allowed to work in. My supervisor Konrad Stahl, though I prefer the German term *Doktorvater*, is an inspiring role-model as an academic with his indefatigable intellectual curiosity and his passion for Science and the Arts. I have greatly enjoyed our discussions and collaboration and will continue to do so. I would like to thank Martin Peitz for his ever-astute answers to questions in many fields and for his momentous support in important choices. And I am very grateful to Elu von Thadden for his invaluable academic input and his dedication to the quality of our graduate program. I am extremely happy and proud to have had the opportunity to work together with Leonardo Felli from whom I have learned much about how to approach problems as an economist. I would further like to thank Annamaria Lusardi for giving me the wonderful opportunity of working in the beautiful and inspiring surroundings of Dartmouth College in Fall 2009.

I have greatly benefitted from my colleagues, both present and past: My graduate school classmates have been fantastic companions and our weekly lunch has been one of my favorite rituals. Frank Rosar, my iCERN co-founder, was always willing to assist me in seminars, as well as game- and auction-theoretic matters. Finally, I truly enjoyed working together and discussing problems with Isabel Ruhmer, Heiko Karle, Tobias Klein, Michal Kowalik, Christian Lambertz, Paolo Massela, Steffen Reik, Dongsoo Shin, Hannes Ullrich and Frank Wachtler – I hope we will have many opportunities to collaborate in the future. I would also like to thank Christina and Kay Wilke for their assistance in the late stages of my dissertation. Further I am very grateful for the help of two outstanding research assistants, Pascal Busch and André Stenzel.

I gratefully acknowledge financial support of the University Mannheim, the Deutsche Forschungsgemeinschaft DFG and the German Israeli Foundation for

Scientific Research and Development. Part of this work was carried out while enjoying the wonderful hospitality of the Study Center Gerzensee for which I would like to thank the Swiss National Bank. Further, I am indebted to the German Automobile Industry Association (VDA) for their cooperation. The views expressed in this dissertation are solely those of the author and not of these organizations.

Finally, yet foremost, I have to thank my family: My wonderful wife Tabea, who with her never-ending support, brilliance and joy has turned this endeavor into something fulfilling and fun. My parents, who with their generosity, patience and love have been true pillars throughout my life, and on whose economic intuition I was always able to rely, as long as there were neither numbers nor formulas involved. And finally my siblings, who were never short of encouragement, and who on some level were happy that one of us went for his PhD – and on all levels were happy it was not them.

Contents

List of Figures

List of Tables

1. General Introduction

Why do achievements differ so widely from aspirations? [...] Social life is not only a trial of strength between opposing groups: it is action within a more or less resilient or brittle framework of institutions and traditions, and it creates – apart from any conscious counter-action – many unforeseen reactions in this framework, some of them perhaps even unforeseeable. To try to analyse these reactions and to foresee them as far as possible is, I believe, the main task of the social sciences. - Karl Popper, *The Open Society and Its Enemies.*

In the course of the research for my dissertation, I had the luck to encounter three fascinating examples for which the introductory quote might as well have been specifically tailored. Social and economic behavior takes place within institutions and is strongly affected by social norms and past choices. The resulting complexity of issues often appears daunting. Take for example the currently existing US patent-regime. Since its inception, the United States Constitution provides that Patent Law was created in order to *"promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries"*. Two centuries have passed, and while the Founding Fathers most certainly had tangible matters like new tools or practical methods to be used in the nascent state in mind, in the meantime the same law must govern the treatment of genome sampling or the source code for computer programs while competing with the patent systems of Europe, Japan, China and many others. This adaptation would not have been possible – and some state forcefully that it has failed miserably – without constant tweaks to the system and continuous marginal reform processes. It is the role of economists to accompany such processes of change and to use the tools at our disposal to better understand the underlying forces and to foresee their effects as far as possible – such as, for example, the current role of patent examiners and judges for the innovation behavior and creative efforts inside a country.

This dissertation is concerned with three different topics, each of which is treated in a separate, self-contained article. Their underlying connection is contained in the quote above, as they apply the methods of social science – microeconomic theory and micro-econometrics – to questions related to "grown", or in one case "growing" institutions. In Chapter 2, I address troubles that patent systems are facing to a large and increasing degree, such as the devaluation of patents due to badly specified claims and property rights and the ensuing central role that courts play in the patent system in the context of a game-theoretic model. While the patent system can be considered a venerable institution, in Chapter 3 I turn my attention to a young and budding one: Online social networks. Making use of a unique data set, in collaboration with Steffen Reik I try to explore whether the existing economic theory on social networks can assist us in understanding the forces that shape user behavior within. One of the central findings is the importance of social norms in this environment. In Chapter 4, which was developed working together with Leonardo Felli and Konrad Stahl, the aim is to show how a specific social norm, *trust*, shapes the interaction between upstream suppliers and downstream manufacturers in the automotive industry. I present a game-theoretic model and test its predictions using data from an extensive industry survey we conducted.

In the following, I briefly introduce the three respective articles which compose the remainder of this dissertation. The Appendix includes most of the tables referred to in the text as well as the complete bibliography.

1.1. Private Profits and Public Benefits - How not to Reform the Patent System

The need for a reform of the US patent system has been widely expressed, both by scientists and practitioners. There are widespread complaints about the quality of granted patents, especially the lack of "clearly drawn property lines", to the extent that some refer to the patent system as "failed". At the same time, due to the investments made into existing patents and the resulting rights of patentees, a fundamental reform of the patent system is not feasible in the short- or medium term. For the foreseeable future, therefore, the system will have to plod on with relatively minor tweaks; but due to its complexity, even apparently minor tweaks can have quite serious consequences.

In order to better understand the tradeoffs involved in the contemplated "improvements", I present a model that includes the three pillars of current patent systems: legislation, patent and trademark offices (PTOs) and courts. Legislators decide the objective criteria concerning the matter that can be patented, such as novelty or utility. The PTO inspects patent applications in compliance with these regulations. The courts, whose importance has increased substantially in the past two decades in the area of patenting, then decide whether claims granted in existing patents are truly enforceable, or whether the patent should perhaps have not been granted initially.

In contrast to the existing literature on patent litigation starting with Meurer (1989), I introduce two different types of lawsuits: Both claims for damages by the patentee against the alleged infringer as well as challenges to the validity of a patent by competitors. I show that it depends on the competitive setting as well as the expected quality of a patent whether litigation can have a complementary function to the examination efforts of the PTO. Against this background, I then show that some of the contemplated (or demanded) reform steps, especially increasing the fees for patenting, may actually decrease the average quality of patent applications in relevant cases.

1.2. Individual (Ir)rationality? Behavior in an Emerging Online Social Network

The last decade has seen important advances in the theoretical literature on the economics of social networks, spearheaded by the work of Jackson and Wolinsky (1996). This article, and its successors, depict social networks using a graph-theoretic environment in which individuals are depicted as nodes and their links – for example business partnerships or friendships – are resembled by the arcs of the graph. The existing theory comes to very precise and relatively homogenous predictions concerning structures of networks that are stable, in the sense that they are the results of equilibrium behavior. But more recently, experimental economists have started to examine the quality of predictions derived from the theory in highly controlled surroundings and the results have been mixed at best, so that doubts have been raised whether individuals truly behave in a rational manner in social-network settings.

In this chapter, I take an empirical look at individual behavior within an actually existing social online-network with additional utility generating functions. To be specific, users are able to upload music to online-libraries, from which both they and their first-degree "friends" can access the music whenever they are online. Standard economic theory proves to be very helpful in predicting with whom individual users are going to form new friendships, as this behavior is compatible with utility maximization. But we also find behavior that points in the direction of social norms that are mutually enforced: Users decline to enter into or sever purely beneficial links when there is indication of free riding, for example. Further, we show that in the face of public-good like provision and organization of music, users do slightly undersupply these forms of effort for higher *levels* of other users' efforts, as standard public-good or club-good theory would predict. But users immediately react to other users' *additional provision* of public goods in a positively reciprocal manner to a degree that more than compensates for the first effect.

1.3. Trust and Investment - A Theoretical and Empirical Assessment

In the fourth chapter, I present a simple model of contractual structures involving upstream suppliers and downstream producers. The upstream suppliers exert effort to determine the joint surplus from cooperating with the downstream firm. Since the latter has the entire bargaining power, this is a classical holdup situation. Defining trust as the expectation of the supplier that the downstream firm will honor his property rights, it is straightforward to show that higher levels of trust will alleviate the typical underinvestment problem. What is more surprising is the relationship between trust and competition among suppliers, which the downstream firm is able to induce. As opposed to the literature on relational contracting, which generally finds that arms-length (market) interactions and functioning informal relationships are to a great extent mutually exclusive, in this setting more intense competition can be associated with *higher*levels of trust. The intuition behind this result is that competition is one mechanism to extract a part of the supplier's surplus that is utilized instead of exploiting the hold-up situation.

We then take these predictions to the data, a unique dataset collected during a two year survey of the German automotive industry. Each of our observations is a relationship between a supplier and a car manufacturer with regard to one part that is procured. This allows us to focus on the relationship specific nature of trust, while existing empirical and experimental studies mostly treat trust as a characteristic of individuals. We are able to study the effects and determinants of trust in these relationships and show that the suppliers' trust, captured by different measures with individual connotations, can be damaged by rent-extracting behavior of the downstream firm. On the other hand, as predicted by our theoretical model, trust is not negatively associated with stronger competition between suppliers induced by the car manufacturer. Emphasizing the importance of trust in vertical relationships, we then demonstrate that trust can mitigate the underinvestment issue resulting from the hold-up situation:

In particular, lower levels of trust of the supplier in the car manufacturer are associated with higher failure rates of the procured part, which we use as a proxy for investment.

2. Private Profits and Public Benefits - How not to Reform the Patent System

2.1. Introduction

2.1.1. Motivation

A patent is not a property right. Patentees cannot directly exclude others from markets that are covered by the claims of their patent(s). Instead, a patent only conveys to the holder the right to go to court against alleged infringers. In recent years, the number of court cases involving patents has grown extremely rapidly to impressive levels.[1] In a thought-provoking study, Bessen and Meurer (2008b) estimate that litigation costs related to patents have exceeded private profits from patents ever since the late 1990s. According to their estimates, global profits directly connected to US patents in 1999 accrued to 9.3 billion USD, while their estimate for the domestic litigation costs to companies is roughly 16 billion USD. These figures are alarming.

The need for a reform of the US patent system has been widely expressed, both by scientists and practitioners. There are widespread complaints about the quality of granted patents, especially the lack of "clearly drawn property lines", to the extent that some refer to the patent system as "failed". Currently, there are efforts underway in both Europe and the US to improve the way patents are granted and enforced.[2] Researchers' more audacious recommendations for

[1] See Cook (2007) for an analysis of this phenomenon, with a focus of the role of specialized courts.

[2] Recent efforts in Europe include the development towards a unified "European Patent" agreed upon in

this have ranged from advising fundamental, game-changing reform such as Bessen and Meurer (2008a), to demanding the patent system as a whole being scrapped such as Boldrin and Levine (2002, 2008).

Compared to these demands, the actual reforms that have been attempted are extremely tame, focussing on (seemingly) minor procedural measures. And also the bulk of theoretical research has focused on more marginal shifts. The reason for this is fairly obvious - despite the many shortcomings of the current systems, the existing patents held by individuals and corporations are tremendously valuable. The investments that lead to the protected innovations were (arguably) sunk with the goal of patenting in mind, which means that any legislature tampering with the system will face never before seen claims for damages. As the youngest current patents have another 20 years on the clock, any fundamental reform steps are all but certain to be postponed until then, at least.[3]

So for the foreseeable future, any revolution to the patent system by needs will be extremely marginal, infused with apparently harmless tweaks to the patent application fee or the responsible jurisdiction for patent claims; but in fact, due to the immense complexity of the current patent-machinery, pitfalls loom even for tweaks that appear to be harmless. The goal of this article is to provide a theoretical framework that is as simple as possible, while it still spans all three pillars of current patent systems. Our aim is to demonstrate a number of perhaps unexpected tradeoffs involved in tampering with them. What we consider the three pillars are:

(1) Patent legislation, which defines the prerequisites for a patent to be granted as well as the right of patent holders.

the Council of Ministers of the European Union on December 4, 2009, as well as the introduction of a specialized European Patent Court as a part of the (currently stalled) European Patent Litigation Agreement. In the US, introducing new patent legislation has become a bi-annual tradition, of which the Patent Reform Acts of 2009, 2007, 2005 and 2003 bear witness. The main goal of the latest was to reduce the burden of the Patent and Trademark Office in the inspection process.

[3] With this point in time moving along until any first steps are at serious reform are taken. And the 20-year timeframe is actually rather optimistic, because research programs that are currently underway with patents potentially looming years from now may already have generated claims based on the investors' trust in the current system.

(2) Patent and trademark offices (PTOs), which apply the existing rules and regulations in examining patent applications.

(3) Courts, which uphold the rights of patentees against alleged infringers and, by giving third parties the chance to challenge existing patents, have the ability to invalidate patents that should not have been granted.

The basic dilemma in patenting is that the patentability of an idea from the perspective of legal requirements or social welfare does not necessarily coincide with the private benefits an inventor obtains from being granted patent claims. In our model, we distinguish between the profit an individual or firm may derive from an idea if patented and the patentability of an idea from the legal perspective. Unlike the existing literature, which mostly focuses on binary *good vs. bad* distinctions, we allow for a continuum along both dimensions - ideas may range from patentable over "almost" patentable to "clearly not" patentable, as is arguably the case in reality. It is the role of the PTO to make this distinction, with its decision necessarily being imperfect - as PTOs face time, budget and staffing constraints. We represent this in a highly reduced form in the model. Finally, there is the possibility of patent holders and competitors becoming entangled in court proceedings, where we can distinguish between damage suits against alleged infringers on the one hand and challenges of the validity of granted patents on the other. We attempt to represent both kinds of proceedings in our model.

The occurrence, type and outcome of court interactions depend on various factors: The most important are the average quality of patents, as there is asymmetric information about the validity of *individual* patents, and the competitive setting between patentees and potential infringers. We show that one of the most frequently contemplated (or demanded) reform steps, increasing fees for patent applications at the PTO, may actually *decrease* the average quality of patent applications. Further, we point out some unintended and potentially costly consequences of increasing the level of scrutiny exerted by the PTOs. Further, we find that a reduction of litigation costs for patent costs has multiple desirable effects - which is especially relevant in the face of specialized courts of appeals being introduced in the US and the European Patent Legislation Agreement.

The remainder of the article proceeds as follows: After a brief review of the related literature, section 2 sketches legal and empirical differences between the European and the US patent system in order to derive stylized facts and motivate the relevant determinants of our model. Section 3 develops the model of patent legislation, examination and litigation. In section 4 we derive the equilibria of the model and perform comparative statics exercises to determine the effects of various approaches to patent system reform. Section 5 emphasizes empirically relevant predictions, proposes venues for further research and concludes.

2.1.2. Related Literature and Contribution

The first strain of the literature that our article is related to studies the effects of patent litigation. The issue of potentially invalid patents and resulting litigation has garnered recurring interest in the past two decades. Meurer (1989) proposes the first model that explicitly takes the possibility of patent invalidity and resulting litigation into account. Patentees face a single competitor, with both parties aware of the fact that courts may overturn the patent. The competitor can either accept the patentees non-cooperative bargaining offer, do nothing or challenge the patent through the courts. The study carefully shows how settlement and litigation probabilities depend, among other things, on the strength of the patent both under symmetric and asymmetric information. We adapt the model for one of our litigation subgames.

Crampes and Langinier (2002) study the opposite competitive setting: Here, the competitor does not challenge the patent, but simply enters the market, thereby (potentially) infringing the patentee's rights. The patentee does not necessarily notice this, instead he must invest in a costly monitoring technology, which determines the probability with which infringement is observed. If observed, the patentee must make the decision whether to accommodate, litigate or settle with the infringer, where the settlement results from cooperative Nash Bargaining. The otherwise efficient bargaining potentially breaks down as the authors impose costly frictions in the bargaining process. Monitoring is shown to be an effective way to preclude entry. Bessen and Meurer (2006)'s approach is very is similar to Crampes and Langinier (2002) - they give the potential infringer the strategic opportunity to invest into research, which affects the probability of being fined by courts in two possible ways: This activity could be general R&D (e.g. adding additional features to the product), which would potentially increase the risk of being found infringing. Or the company could exert effort to specifically invent around existing patents, which would reduce the probability of infringing. Instead of monitoring activity as in Crampes and Langinier (2002), the patent holder can invest in strengthening his patent; the probability that the court will find the infringer guilty if a trial arises increases in this investment (this can be interpreted as applying for additional patents to create a thicket or employing better and more expensive lawyers to phrase the patent application and claims optimally). In the empirical part of their project, the authors find that competitors' investment is overwhelmingly focused on research.

More specific issues within the area of patent litigation have also been discussed: A number of papers investigate the role of different liability rules and how these affect the level of protection granted by intellectual property rights, see e.g. Schankerman and Scotchmer (2001) and Choi (2006) for comparisons of lost profit and unjust enrichment rules in different competitive settings. Anton and Yao (2007) study the infringement decision in the case of process innovations for lost profits damages. Antitrust issues and anticompetitive effects of the settlement of patent litigation are taken into account in Shapiro (2003) and Lemley and Shapiro (2005).

A second group of articles is concerned with the organization of and incentives within patent offices and their effects on the patent system. Caillaud and Duchene (2006) scrutinize the capacity of the patent office to deal with applications and the so called overload problem. In their model, if the PTO were to employ a strict examination standard, a hypothetical equilibrium exists in which bad applications are deterred. A problem arises, though, if the PTO is overloaded with applications - as it can no longer enforce the strict standard, deterrence becomes less effective and the separating equilibrium can no longer be upheld. In Prady (2008), low quality inventors can induce shirking of the patent examiner by sending signals that require more effort to disentangle. Schuett (2009) models the examination process as a combined moral-hazard/adverse-selection problem and tries to explain and argue for different incentive schemes for examiners.

Two recent studies are closely related to ours. Both Farrell and Shapiro (2008) and Chiou (2008) present models which allow for patent examination and potential litigation. In Farrell and Shapiro (2008), an upstream innovator owns a patent that may be "weak" in the sense that a court will only uphold it if challenged with a certain probability. The upstream innovator faces a set of downstream firms which can apply his innovation. Depending on the competitive setting, they show that even very weak patents can have a strong price-shifting effect – when the downstream market is very competitive, then there is next to no private incentive to challenge the patent, which gives even weak patents great power. In such a setting, more stringent PTO reviews are welfare enhancing. Our study is in many ways complementary to this approach: While in Farrell and Shapiro (2008) there is no room for litigation in equilibrium, despite zero litigation costs, our main focus is to determine under which conditions which kind of litigation will appear, and how this is affected by contemplated reforms. Their main analysis focuses on the role of the intensity of downstream competition in a specific setup - we treat this factor in a reduced form, which to a certain degree encompasses their analysis. Finally, they treat the quality of granted patents as uncertain even for the holders. We demonstrate that similar effects can be generated with asymmetric information and two kinds of patents –ironclad or void – and that the probabilistic nature is no requirement.

Motivated by Lemley (2001)'s provocative thesis that increases in the diligence of the patent office are inefficient and one should let the competition sort out bad patents through litigation, Chiou (2008) sets up a two stage model, in which first the patent office exerts effort in order to find prior art that allows to deny a latently invalid patent. If the patent office is unable to destroy the patent, in the second stage a private competitor can exert effort in the same way to have it repealed. While for relatively good patents (high priors of patent quality) he finds a crowding out of private efforts through the patent office along the lines of Lemley (2001), for low-quality patent population he finds the opposite effect: Stricter enforcement by the patent office actually encourages the competitor to himself invest in research efforts to try to invalidate the patent. The paper relies on a new way of modeling court interactions – higher exertion of (costly) effort by the challenger (or by the PTO) leads to the destruction of the patent with a higher probability and the "quality" of a patent affects the marginal effect of effort. Again, our approach, using a more classical way to model suits and countersuits applied to a population of patent applications, can be seen as complementary.

This article contributes to both strands of the literature. Our model encompasses both the examination and the litigation phase of the patenting process. Regarding litigation, we integrate the competitors' choice between challenging an existing patent or entering the market directly and potentially infringing. Competitors' beliefs concerning patent validity are determined by the policy adopted by the patent office as well as settlement offers proposed by the patent holders. Further, we allow for a delay between patent application and the final decision of the patent office - a form of patent office overload related to the one studied by Caillaud and Duchene (2006), which can be considered a "cost" of demanding higher levels of diligence from patent examiners and which awaits empirical exploration. Finally, we study a population of ideas which varies continuously regarding private value to the inventor and objective patentability, which allows us to shed light on tradeoffs due to the composition of the population of patent applications, which have not been studied before.

2.2. Comparison of the US and the European Patent System - Stylized Facts

Before we describe the theoretical model, we first provide a glimpse of some of the differences between two of the most important patent regimes today: the US and the European patent system. The goal is not to exhaustively display and analyze these issues, but to motivate the factors and variables we integrate into our model.

2.2.1. Patent Legislation and Examination through Patent Offices

Even though there are many calls for and initiatives in the direction of harmonization, the US and the European patent legislations retain some very distinct features. Generally speaking, there are stark differences between European and US patent laws, while the differences among the core countries of the EU are comparably negligible.

While the "first to file" system is prevalent in Europe, that is, the first to apply for a patent is awarded the right, traditionally the US award patent rights to the first person to discover an innovation ("first to invent"). A further difference regarding the patenting requirements concerns what in the US is termed "novelty" and, e.g., in Germany the "inventive step". Both requirements state that for an idea to be patentable it may not yet have been made public previously in any form. In Germany this holds absolutely. In the US, inventors are granted a "grace period", i.e. they generally only have to file their application within a year of their idea's publication.

The other general requirements for something to be patentable are similar in both systems - patentable subject matter, non-obviousness, and applicability. Still, there is quite a lot of evidence that these requirements are interpreted differently by the US and European patent offices, which is close to the focus of our study. For example, Straus and Klunker (2007) cite the following numbers of the US and the European patent offices: In 2005, there were 409,532 applications for patents in the US. Out of these, 165,485 were granted, i.e. a share of 40.4 per cent. On the other hand, in Europe, the number of applications in the same year was 197,391 with 53,256 patents being granted, i.e. a significantly smaller share of 27,0 per cent. From these numbers alone, it appears that that the German patent offices apply a stricter standard than the American ones. This is corroborated by other observations: The average time from filing a patent to it being granted was 45,3 months in Europe, while it was only 24 months in the US, as Hall and Harhoff (2004) report. They also look at the grant-rates of patents at the European Patent office for US patents seeking European approval on the one and patents from other countries seeking approval in Europe on the other hand. They find that the approval rate for US patents was substantially lower, with a 16 per cent difference in 1995. The trend points towards a deterioration of the US standards, as in 1979 there was parity concerning these numbers.

Straus and Klunker (2007) further cite a study by the consulting firm Roland Berger that found that in order to apply for and maintain a patent over its entire lifetime, costs between 32,000 and 47,000 EUR accrue in Europe, while the average figure is only 10,250 EUR in the US.[4]

From these observations, we derive the following stylized facts regarding the patent application process in the US and in Germany:

(F1) Patent applications are scrutinized more strictly by European than by US patent offices.

(F1') As a result, a European patent is potentially a better indicator of the strength of the patentee's claims than a US patent.

(F2) It is significantly more expensive to obtain a patent in Europe than in the US in terms of the fees required.

[4]For comparison: Regular period from filing to grant: 31,6 months in Japan, 45,3 months in Europe, 24 months in China and US. Patent costs are estimated to range from about 2,400 to 4,000 EUR in China.

(F3) The duration of patent pending varies between patenting systems by a considerable margin.

2.2.2. Courts and Patent Trials

After briefly discussing the differences concerning the laws in the two jurisdictions, we will next consider how patent suits fare in the courts.

Cook (2007) notes that in the course of the last decade, the number of patent cases filed in the US has roughly doubled, from about 1,250 in 1990 to about 2,500 in 2000. He shows that this is one the one hand a result of more patents being filed in general (the ratio of cases per granted patent remains relatively unchanged, as the latter number increases from about 90,000 to about 180,000 per year in the same period). On the other hand it results from courts being more accommodating to plaintiffs. In his empirical model the decision to go to court depends, among other factors, on the share of cases that were successful in the respective district in previous periods. One impressive measure in this regard is that the probability of receiving a patent reward of 1 million USD or more in 2001 constant dollars increased from less than 10 percent in 1976 to more than 30 percent in 2000.

Concerning the average costs of going to court over patent infringement in the US, estimates range from 500,000 USD to 3,000,000 USD for each party. Much larger sums are mentioned in the context of complex or high-stakes cases, especially in the area of pharmaceuticals. For a European patent to be litigated, court fees of 70,000 EUR arise according to European Patent Office (EPO) reports. Their estimates of the lawyers' fees bourn by the parties in addition amounts to approximately the same figure, so that each party would have to expect costs around 150.000 EUR ex ante.

We glean the following stylized fact from these brief observations:

(F4) There is a significant cost difference between bringing a case to court in the US and Europe, i.e. the costs of litigation are vastly higher in the US than in Europe.

2.3. Model Setting

2.3.1. The rules of the patent system

Legislators (exogenously) define a minimum standard for utility, novelty, non-obviousness, etc which makes an idea objectively patentable. For this model, we assume that these various criteria can be reduced to one dimension, and the minimum objective standard set by legislators is denoted μ.

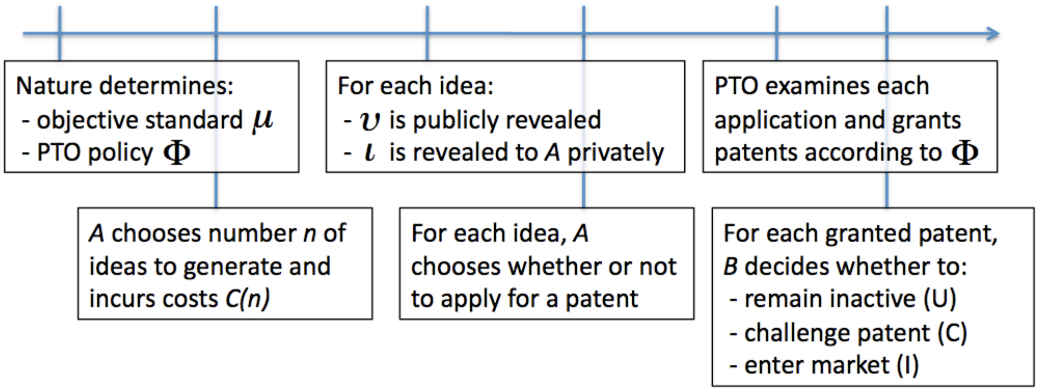

Figure 2.1.: Timeline of the entire model

2.3.2. Strategic Players and "Ideas"

We consider a game with two risk-neutral strategic players, A and B. Player A actively generates ideas. Ideas are "random events", i.e. individual ideas are imbued with certain randomly drawn characteristics (more on this below). Generating n ideas in a given period, A incurs costs $C(n)$. We assume that $\frac{\partial C(n)}{\partial n} > 0$ and $\frac{\partial^2 C(n)}{\partial n^2} > 0$.[5]

Each individual idea that is generated is defined by the following characteristics:

[5]This can be interpreted in the following way: A resembles the population of potential inventors and they are sorted according to the marginal costs of their ideas.

(1) The patentability of the idea from a legal standpoint, i.e. its utility, novelty, non-obviousness etc. Again we assume that these can be reduced to a single scalar, which we denote by ι, with $\iota \epsilon [0, 1]$. Note that if $\iota \geq \mu$ an idea is objectively patentable from the legal standpoint.

(2) A measure of the additional value to the inventor derived from patenting the idea, which we denote by υ. To be more specific, this measure is defined as the value that the inventor can appropriate from exploiting the idea optimally as a monopolist after obtaining the patent, i.e. given no-one infringes upon the patent. As we are concerned with the incentive effects of patents in the face of challenges with respect to enforcement, we set the value from exploiting an idea without patent protection to zero. Therefore in our model, from an ex ante perspective, an idea is only valuable to its inventor in as far it can be patented. We normalize once more such that $\upsilon \in [0, 1]$.

The inventor only learns the actual characteristics of a given idea *after* it has been generated and he has incurred the associated costs. Specifically, the parameters ι and υ are respectively drawn from commonly known and independent[6] distributions $F(\iota)$ and $G(\upsilon)$ and revealed to A. We further assume that F and G are both continuous and strictly increasing.

2.3.3. Patent and Trademark Office

Patent Applications and the Patent Office

A decides for which of his ideas to submit patent applications. It costs A the fixed amount τ_1 to submit an idea to be examined by the patent office. We assume that the PTO can only make a binary decision regarding a given patent application: either approve or reject it. According to the legal rules stated above, the PTO should approve patent applications if and only if $\iota \geq \mu$.

[6] We only use this assumption (independence of distributions) to be able to derive a closed form solution for an expression below. The majority of our results hold without it.

We assume that the PTO cannot directly observe the actual ι of an individual application. Conceptually therefore, it could commit two kinds of errors: Errors of type one, i.e. granting a patent despite $\iota < \mu$, and errors of type two, i.e. declining patents despite $\iota \geq \mu$ and both of these errors are associated with social costs. For the sake of this model, we neglect errors of type two - discussions both with scientists and practitioners have convinced us that they are empirically close to irrelevant, since there is a relatively fast and cheap way to appeal the decision in both jurisdictions in this case. Imagine that A can incur the additional (small) costs τ_2 in order to have the application reexamined which leads to it being granted if $\iota \geq \mu$. The only effect of this will be to shift the expected profits from patenting of holders of patentable ideas downwards by a fixed sum.

We model the patent office policy in the most simple possible reduced form. The patent office simply implements an exogenously given examination policy Φ which determines the likelihood $\phi(\iota)$ of a given idea being patented, including the probability of a first order mistake being made if $\iota < \mu$. For example, if the PTO grants every patent application, then $\phi(\iota) = 1$ for any ι. The inspection policy does not depend on υ – for one, patent examiners (as everyone else) have a very hard time determining the expected value of a patent, and also, more importantly, guidelines in general forbid different treatment of applications depending on their suspected value. We assume that $\phi(\iota) > 0$ for all ι, i.e. even the worst quality idea always has a positive chance to be awarded a patent. Further we assume that $\frac{\partial \phi}{\partial \iota} > 0$ if $\iota < \mu$, i.e. the closer ι comes to the objective patenting threshold, the more likely a patent is to be (falsely) granted. Therefore the probability with which a given application will be granted is $\phi(\iota) \leq 1$ if $\iota < \mu$, and $\phi(\iota) = 1$ if the idea is objectively patentable.

[Consider the following example for a possible micro-foundation as an illustration of the idea: The patent office generates a number ζ of signals, where ζ is part of the exogenous schedule Φ. ζ can for example be interpreted as the time that a patent examiner spends researching prior art or perusing the patent application. If individual signals are normally distributed with mean ι

and standard deviation σ*, the mean of the signals* $\bar{\iota}$ *has the expected value* ι *and the standard deviation* $\varsigma = \frac{\sigma}{\zeta^{0.5}}$*. The examination schedule further includes a cutoff level* μ^**. If the mean of the generated signals is below this cutoff level, the PTO rejects the application, and it grants the patent otherwise.* ζ *and* μ^* *therefore implicitly determine the probability of type one and type two errors that the PTO commits for any given* ι*. Denote the normal distribution with mean* ι *and variance* ς *as* $N_{\iota,\varsigma}$*. Then the probability for a type 1-error given that* $\iota \geq \mu$ *is simply* $N_{\iota,\varsigma}(\mu^*)$*, and equivalently the probability for a type 2-error (which we do not consider) given* $\iota < \mu$ *would be* $1 - N_{\iota,\varsigma}(\mu^*)$*.]*

Extension: Period of Patent Pending

As discussed above, the decision whether or not a patent application will be granted is by no means made immediately. Instead, a substantial period of time passes between application and decision, during which the ideas has the status of "patent pending". Little investigation has focused on the value of patents in the course of the application process so far. But when Steve Jobs exclaimed his now famous "And, boy, have we patented it!" during the speech introducing Apple's iPhone at the MacWorld Expo in January 2007, the greatest share of the patents protecting the touch-screen technology involved *had not been granted, yet.* In fact each of the 21 iPhone-patents considered "central" by technology-afficionadoes was still pending at this point in time.[7] Part of this stance was surely justified by the knowledge that the patents would be granted, later on, which turned out to be true. But to some extent, pending patent applications itself are valuable, especially in fast-moving industries.

There are a number of reasons to suspect that inventors do benefit from their ideas during the pending period. Many if not most license agreements are negotiated prior to the actual patent being issued. Other firms may think twice about entering the market if the product is designated with the "patent pending" stamp, fearing future lawsuits. Finally, in one of the few empirical

[7] See http://www.mad4mobilephones.com/the-21-most-important-iphone-patents/562/ for an extremely detailed description and links to the individual claims.

studies related to the topic, Häussler et al. (2009) find that the chance to receive venture capital rises with the number of patent applications in the firms' portfolios.

Again, we reduce this concept as far as possible. We simply denote the patent-pending value of an idea (ι, υ) as $\delta(\Phi)\upsilon$, where $0 \leq \delta \; < 1$ and $\frac{\partial \delta}{\partial \phi} < 0$. Therefore increasing the quality of patent office scrutiny (or lowering $\phi(\iota)$ is "expensive" in the sense that it increases the patent pending value of ideas, for example by increasing the duration between application and examination on average. We will not over-stretch this concept in the following, especially in this overly simple linear specification, but it is useful to remember this factor when discussing reforms aimed at more stringent inspection policies through the PTO.

[Consider again the illustrative example above. Imagine that the patent pending duration of a given idea after application is t *periods, where* $t = t(\zeta, n)$ *and* $\frac{\partial t}{\partial \zeta} > 0$*. During this time, the perspective patentee receives the benefits* $\Delta(\upsilon, \iota)$ *per period.]*

2.3.4. The Court System

We have described above how ideas are generated, how A decides whether or not to apply for patents, and whether or not patents are granted by the PTO. In the next and final stage of the model, B enters the picture. We consider B to be the only strategically acting competitor of A in each of the markets covered by a patent in the economy.[8] B observes which patents were granted to A and chooses one of the following reactions: He may either leave the corresponding market uncontested (U), challenge the validity of the patent before court (C) or enter the market and thereby infringe A's patent (I).

Choices (C) and (I) respectively induce different litigation games: For (C),

[8] Analogously to above, B can be interpreted as the population of potential competitors to the population of inventors A. More specifically, B is the most efficient or profitable competitor of A - this mitigates the loss of generality from assuming only one competitor, as other less efficient firms' incentives are aligned with B's, yet less strong.

we adapt the approach introduced by Meurer (1989). For choice (I), we use a setup related to Crampes and Langinier (2002), but with non-cooperative bargaining.

Information

While the patent holder is perfectly informed about the quality (ι, υ) of any given patent, B cannot observe ι. This parameter has the following significance in the context of court proceedings: The probability that a court upholds a given patent obviously depends on the legal patentability of the underlying idea. For simplicity, we assume that courts make the "right" decision, that is, they uphold the validity of the patent whenever $\iota \geq \mu$. One might question the assumption that courts can discover the true ι of an idea, while the PTO cannot. But seeing that court cases take years while patent examiners have days (if that!) to come to a decision regarding a given application, the assumption may appear less severe.

Therefore, from A's perspective, there is no uncertainty regarding the court's decision on the validity of a given patent: if $\iota \geq \mu$ the court will uphold its claims, and it will repeal the patent otherwise. As B cannot observe ι, though, the court proceedings take on a random character from his point of view and he must form the subjective probability of a patent being upheld given the information available to him. In our simple model, B only observes three things: First, the fact that a patent was granted by the PTO. Second, the value υ of the patent to the patent holder, which he may e.g. derive from observed sales. Finally third, eventual settlement offers S proposed in the course of the litigation game by A. From these observations, B forms the conditional subjective expectation $p_B(\iota \geq \mu | \upsilon, \Phi, S)$ of the validity of the patent. We will abbreviate the probability unconditional on a settlement offer S as p_B in the following and denote updated equilibrium beliefs following an informative signal S as p_B^*.

Structure of the Post-Patenting Litigation-Game

Combining the ingredients above, as the final step of the model, for each patent that was granted to A, the following subgame is played:

B initially decides whether to leave the patent uncontested (U), challenge the validity before a court of law (C) or enter the market and infringe A's patent (I).

(U) If B leaves the market uncontested, the game regarding the given patent ends and the players get the payoffs $\pi_U^A = \upsilon$ and $\pi_U^B = 0$.

(C) This case is treated analogously to, for example, Meurer (1989) and Chiou (2008). If B announces her intention to challenge the patent, A and B first have the opportunity to negotiate a settlement. The informed party, i.e. A, who knows whether or not the patent will be held valid by the court, can first propose a settlement offer S_C. B observes the settlement offer and in this process updates her subjective probability that the patent is valid. Then she decides whether or not to go to court, with the following outcomes: If $\iota \geq \mu$ the court upholds the patent despite the challenge and the payoffs are $\upsilon - \kappa$ for A and $-\kappa$ for B. If $\iota < \mu$ the court rules that the patent is invalid and therefore the market is no longer protected. By definition, the payoff of A therefore is $-\kappa$ and B obtains the payoff from a competitive market with free entry net of court fees which we denote as $\upsilon_B^C - \kappa$.[9]

(I) The following subgame ensues if B unilaterally enters the market with a product that potentially infringes A's patent. We assume that (C) and (I) are mutually exclusive, i.e. it is impossible to, for example, challenge a patent *and* enter the market.[10] Again, we assume that it is the informed

[9]Note that any market payoff always depends on the overall profitability of the market, therefore for example υ_B^C is actually $\upsilon_B^C(\upsilon)$. We continue to use the former as a shorthand.

[10]There are a number of sound economic reasons underlying this assumption. For one, challenging a patent before introducing a product raises the specter of willful infringement and punitively higher compensation in the case of being found infringing - we leave this consideration out of the model. Further, if the company

party A who makes the settlement offer S_I. If B rejects the offer, A decides whether or not to litigate. As the patent is either valid ($\iota \geq \mu$) or not valid and A is perfectly informed about this, the case in which he sues and the courts invalidate the patent does not arise (here, A would receive a payoff of $-\kappa$ and B receives the competitive payoff net of court fees $\upsilon_B^C - \kappa$; therefore it is better for A to accommodate given patent invalidity, which gives the duopoly payoff to both parties). If the patent is valid and A sues, the courts grant him lost profits or compensatory damages, therefore his payoffs from litigation are $\upsilon - \kappa$.[11] On the other hand, the payoff of the infringer in this case is her duopoly profit plus the difference between A's duopoly profit and the forgone monopoly profit minus court costs, i.e. $\upsilon_B^D - (\upsilon - \upsilon_A^D) - \kappa$.

The Structure of Industry Payoffs

We make the following assumptions with regard to the industry profits. First, we assume that $\upsilon > \upsilon_A^D + \upsilon_B^D$, the profit that A can optimally extract is larger than the sum of the duopoly profits of A and B.[12] Second, we assume that $\upsilon_B^D \geq \upsilon_B^C > 0$. B makes weakly higher profits in duopoly than under full entry competition, but still makes strictly positive profits in the latter case (though they may be arbitrarily close to 0). Third, we assume that υ_A^D, υ_B^D and υ_B^C are all continuous and increasing in υ - loosely speaking, a *more* valuable market in general leads to higher duopoly and competitive payoffs for market participants. Each of these would hold for generic specifications of Cournot competition with entry under integer constraints, for example. Fourth and

is not found to be infringing, it may itself benefit indirectly from the patent protection of the market, as we will see below.

[11] We only focus on this form of damages, for an excellent discussion of the effects of different kinds of damage awards see for example Schankerman and Scotchmer (2001) or Choi (2006). Further, we abstract from the fact that an alleged infringement of a valid patent is not necessarily covered by the claims of the patent. One could simply rescale the payoffs of the players by the ex ante probability that the product will be found infringing, without significant changes to the results. To reduce the notational burden, we abstain from this exercize.

[12] Implicitly, we thereby disregard such phenomena as "patent trolls", companies which cannot exploit their patents effectively *without* going to court (or threatening to do so) against producers who inadvertently infringed upon their patent. See Reitzig et al. (2007) for a study focusing on this topic.

finally, we assume that $(\upsilon - \upsilon_B^C)$ and $(\upsilon - \upsilon_A^D)$ are nondecreasing in υ. As the market as a whole becomes more valuable, the differences between monopoly- and duopoly-, or monopoly- and competitive profits do not grow smaller.

2.4. Analysis

2.4.1. Equilibria

In the following, we first derive the Bayes-Nash Equilibria for the litigation subgames, before embedding them in the larger model in order to learn more about the incentives to generate ideas and patent them.

Litigation Subgames

We first focus on the case that has already been studied in the literature, i.e. **subgame (C)** in which B threatens to challenge A's patent before court.[13] Clearly this threat is not credible unless the unconditional expected payoff from going to court is non-negative, i.e. the following condition holds.

$$(1 - p_B)\upsilon_B^c - \kappa \geq 0 \tag{2.1}$$

If (1) is violated, neither holders of valid nor of invalid patents make a settlement offer to B, who in turn does not go to court over the issue. This resembles a pooling equilibrium with the payoffs $\pi_A = \upsilon$ and $\pi_B = 0$.

The clearly more interesting case is the one in which (1) is satisfied. Note that if $\iota \geq \mu$, A cannot credibly signal to B that he is the owner of a valid patent via his settlement offer and keep him from going to court completely, as signals are costless and can be mimicked by holders of invalid patents. Intuitively, no matter which sum holders of good patents offer to their competitor, they can

[13]As the derived results for this subgame have been established previously, we cover this case very briefly. For a detailed discussion of the technical aspects of the equilibrium, we refer to Meurer (1989).

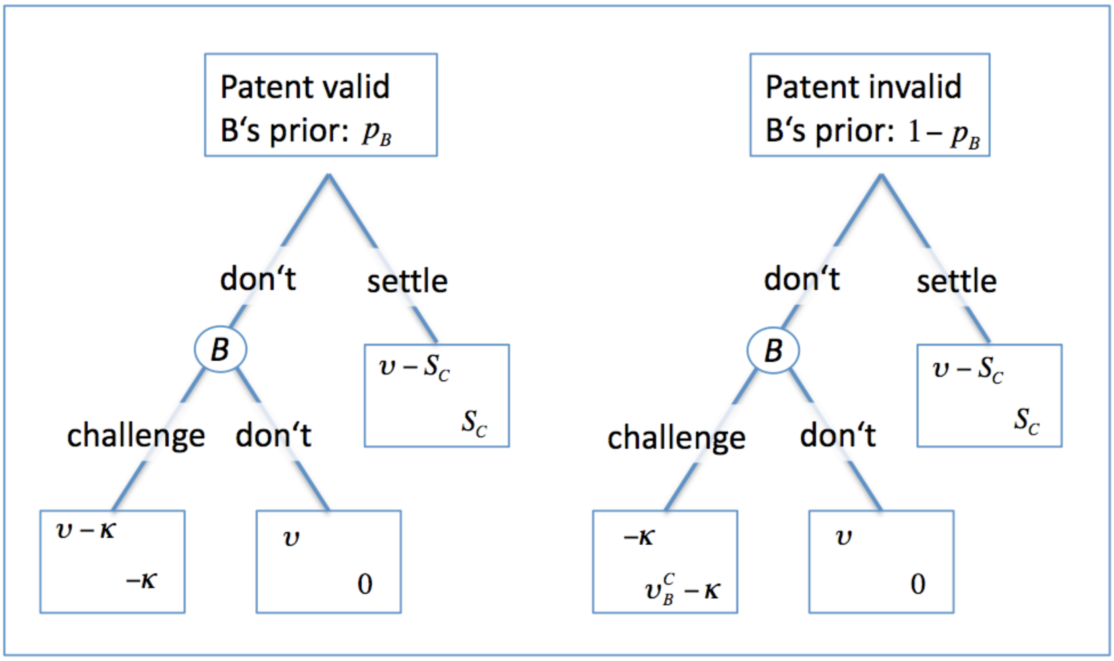

Figure 2.2.: Game Tree for the Challenge Litigation Subgame

always be mimicked by the holders of bad patents and litigation will occur with positive probability in equilibrium as a result. The following is an equilibrium: Holders of good patents make the toughest settlement offer possible, i.e. $\hat{S}_C = 0$. Holders of bad patents mix between $\hat{S}_C$ and offering a settlement worth $\underline{S}_C = \upsilon_B^C - \kappa$.

Lemma 1: *[Proposition 1 - Meurer (1989)] Perfect Bayesian Equilibrium of subgame (C): (i) If (1) is violated, in equilibrium no settlement offer is made and no patent is challenged. (ii) If (1) is satisfied, a semi-separating equilibrium arises in which holders of good patents never make a settlement offer, holders of bad patents mix between making no settlement offer and offering* $S_C = \upsilon_B^C - \kappa$ *and B updates his beliefs to* p_B^* *upon receiving or not receiving a settlement offer and mixes between litigation and inaction.*

We briefly sketch the existence proof in the following in order to clearly demonstrate the mechanics and intuition of the equilibrium. With the refinements of sequential equilibrium and D1, this equilibrium is also unique - for these proofs we refer the reader to Meurer (1989).

Proof. (i) As the expected payoff from litigation given p_B is negative and no signal is forthcoming, no updating takes place and B chooses inaction. (ii) Denote the probability with which B goes to court given the signal $S_C = 0$ as λ. Then the expected payoff of low quality patentees from mimicking the behavior of high quality patentees is: $(1-\lambda)\upsilon+\lambda(-\kappa)$, while the certain payoff given the settlement offer is $\upsilon - (\upsilon_B^C - \kappa)$. For low quality As to be willing to mix between these two, they have to be identical. Consider next the payoff of B if she observes the signal $S_C = 0$. Let us call the probability with which low-quality patentees mimic good quality patentees β. Upon perceiving the signal $S_C = 0$, Bayesian updating of the prior gives us the following condition: $1-p_B^* = \frac{\beta(1-p_B)}{p_B+\beta(1-p_B)}$. For B to be willing to mix between litigation and accepting the tough settlement offer, the following equality has to hold: $p_B^*(-\kappa) + (1 - p_B^*)(\upsilon_B^C - \kappa) = 0$. For the following values of λ and β, all of these conditions are fulfilled simultaneously: $\lambda^* = \frac{\upsilon_B^C-\kappa}{\upsilon+\kappa}$ and $\beta^* = \frac{p_B}{1-p_B}\frac{\kappa}{\upsilon_B^C-\kappa}$. Finally, specify off the equilibrium path beliefs $p_B(\iota \geq \mu|\upsilon, \Phi, S > 0) = 0$. Then it is easy to show that no profitable deviation exists for either player.

Note that litigation arises here as a result of the competitor not being able to distinguish between good and bad patents given that the offered settlement is 0, while she can identify those patent holders as bad who offer a more generous settlement. For these, though, B needs not litigate, as she receives the same expected payoff from the settlement as she would from litigation. Note that the holders of bad patents are exactly indifferent between the tough settlement offer of 0 and the generous offer due to the chance that the competitor will take the case to court with a positive probability given the tough offer, which results in them losing their patent protection entirely.

Let us next consider the **subgame (I)**, in which B enters the market, thereby infringing the potentially valid patent of A. From the payoff-structure above it is clear that it is never profitable for holders of bad patents to go to court, as the court is certain to deem their patent invalid and clearly $\upsilon_D^A > -\kappa$. Holders of

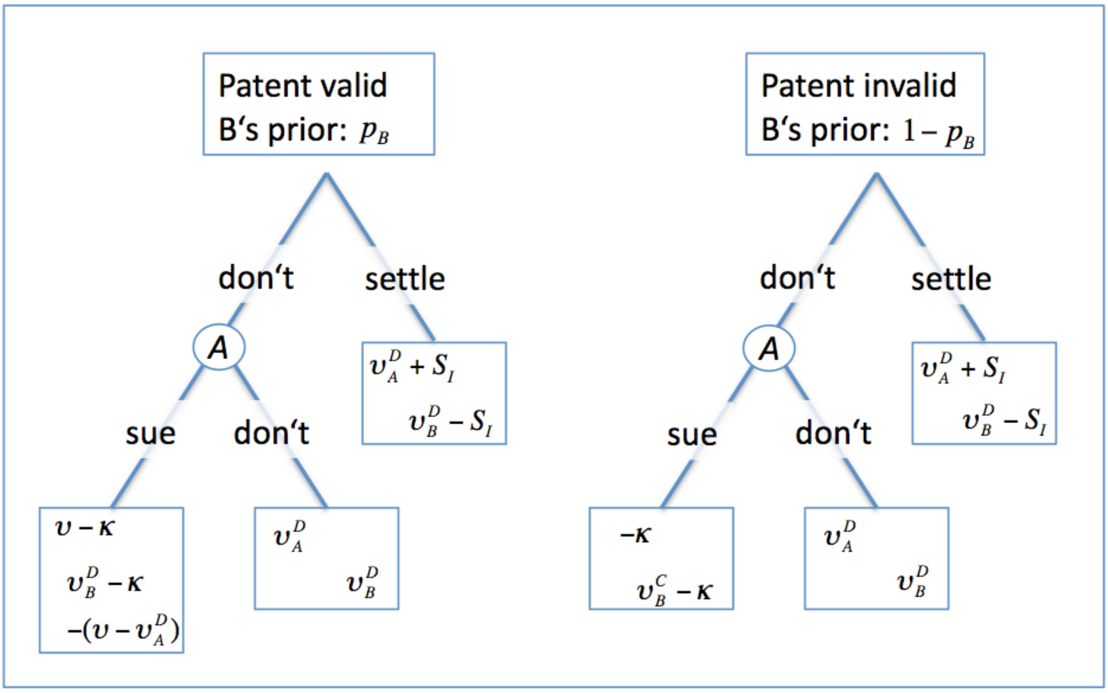

Figure 2.3.: Game Tree for the Infringement Litigation Subgame

valid patents will prefer litigation to inaction only if the (certain) court outcome of monopoly profits net of court costs is larger than the duopoly outcome they would receive given inaction, i.e. the following condition holds:

$$\upsilon - \upsilon_A^D \geq \kappa \tag{2.2}$$

If (2) is violated, even the patentee with a valid patent cannot credibly threaten to go to court, therefore B will not be willing to pay any compensation in settlement negotiations and $S_I = 0$. The payoffs of the two parties are $\pi_I^A = \upsilon_D^A$ and $\pi_I^B = \upsilon_D^B$. Again, the more interesting case is when (2) is satisfied. Here, it is profitable for holders of good patents to sue for infringement damages, while it is still not in the interest of holders of bad patents to do so. But as opposed to subgame (C), now it is the *informed* party who decides whether or not to go to court. Unlike in the previous case therefore there is no positive probability of litigation to induce patentees of the bad type not to mimic the good type and even a semi-separating equilibrium cannot be supported. Therefore in equilibrium, no informative settlement offer exists, in the sense that it allows B to update her prior probability of patent validity.

Good patentees know that they will receive a gain of $(\upsilon - \kappa) - \upsilon_A^D$ from litigation over inaction. The expected loss of B from failing to reach a settlement, on the other hand, is equal to the expected damages payment plus court costs, i.e. $p_B(\upsilon - \upsilon_A^D + \kappa)$. From this we derive the following condition:

$$p_B(\upsilon - \upsilon_A^D + \kappa) \geq (\upsilon - \kappa) - \upsilon_A^D \tag{2.3}$$

If this condition is violated, then holders of valid patents and the competitor are unable to obtain a settlement and the patent holder sues for damages. These considerations allow us to formulate the following lemma:

Lemma 2: *Perfect Bayesian Equilibrium of subgame (I): (i) if (2) is violated, no settlement is reached and A remains inactive in equilibrium. (ii) If (2) and (3) are satisfied , settlement is reached with certainty. (iii) If (2) is satisfied and (3) is violated, no settlement is reached. Upon failure, holders of valid patents sue and holders of invalid patents remain inactive.*

Proof: (i) When (2) is violated, the threat of litigation is not credible for either kind of patentee, therefore no non-negative settlement demand will be met. (ii) Define a settlement offer $S_I^* \epsilon [(\upsilon - \kappa) - \upsilon_A^D, p_B(\upsilon - \upsilon_A^D + \kappa)]$. Then the following pair of strategies is a Perfect Bayesian equilibrium: Both types of A offer S_I^*. B accepts S_I^* and declines any other settlement offer. If a settlement offer has been declined, holders of valid patents sue and holders of invalid patents remain inactive.[14] (iii) Denote good quality A's settlement offer as S_I^+ and bad quality A's offer as S_I^-, with $S_I \geq 0$. We know that good quality patentees will demand no less than their certain gains from litigation, so that $S_I^+ \geq (\upsilon - \kappa) - \upsilon_A^D$. Now assume that $S_I^+ > S_I^-$ and B accepts S_I^+

[14]Clearly, there is a continuum of equilibria. In the following, as A can make a take it or leave it offer, we select the equilibrium most profitable to A whenever we consider payoffs from the subgames.

with positive probability. Then bad quality patentees will want to mimic and deviate to S_I^+. The same holds for the opposite case, therefore $S_I^+ = S_I^-$ in equilibrium. As the minimal settlement offer of good inventors is larger than the expected court outcome for B, i.e. $S_I = (\upsilon - \kappa) - \upsilon_A^D > p_B(\upsilon - \upsilon_A^D + \kappa)$, no feasible settlement offer exists that B is willing to accept.

It is convenient to summarize the results of Lemma 1 and 2 in the table below.

	Outcomes	$\pi_A(\iota \geq \mu)$	$\pi_A(\iota < \mu)$	π_B
Case (C)				
(1) violated	no S, no L	υ	*	0
(1) satisfied	*all* (mixed)	$\upsilon - \frac{\upsilon_B^C - \kappa}{\upsilon + \kappa}\kappa$	$\upsilon - (\upsilon_B^C - \kappa)$	$(1 - p_B)\upsilon_B^C - \kappa$
Case (I)				
(2) violated	no S, no L	υ_A^D	*	υ_B^D
(2) sat, (3) vio	L if valid	$\upsilon - \kappa$	υ_A^D	$\upsilon_B^D - p_B(\upsilon - \upsilon_A^D + \kappa)$
(2) & (3) sat	S	$\upsilon_A^D + p_B(\upsilon - \upsilon_A^D + \kappa)$	*	$\upsilon_B^D - p_B(\upsilon - \upsilon_A^D + \kappa)$

Table 2.1.: Equilibrium outcomes and expected profits of the litigation subgames
(S indicates settlement, L litigation. When holders of invalid and valid patents receive the same profits, these are abbreviated by an asterisk.)

In passing, note that we get the familiar result in the challenge subgame (C) that weaker patents are more likely to result in the settlement outcome. More interestingly, the infringement subgame (I) delivers the opposite result - here, if the perception of patent quality is highest, settlement occurs (both for valid and invalid patents), while intermediate patent quality lets litigation arise (only for valid patents). An analysis focusing on only one type of potential litigation will necessarily overlook this fact.

The combination of Lemma 1 and 2 allows us to derive our first proposition:

Proposition 1: *In the Perfect Bayesian Equilibrium of the litigation subgame, B will prefer challenging an existing patent to infringement and inaction iff the following hold: (i) Condition (1) is satisfied, (ii) condition (2) is satisfied and (iii)* $p_B > \frac{\upsilon_B^D - \upsilon_B^C + \kappa}{\upsilon - \upsilon_A^D - \upsilon_B^C + \kappa}$.

The proof of the proposition follows directly from the preceding Lemma 1 and 2. Briefly note that if (1) is violated, B is indifferent between alternative (C) and inaction, and when (2) is violated B always prefers infringement, as this gives him the duopoly payoff. Condition (iii) of the proposition is perhaps the most surprising: This shows that only for patent quality priors *above* a certain threshold B prefers the challenge subgame over infringement. Intuitively, as the patent quality increases, the higher likelihood of having to pay damages makes infringement relatively less attractive.

This lower threshold increases (i.e. challenges become less likely) as the costs of going to court κ increase - this favors infringement, where it is the patent holder who must decide whether or not to initiate court proceedings. It further increases in the difference between duopoly and competitive profits for B, $\upsilon_B^D - \upsilon_B^C$. The higher the protection that B enjoys passively from A's patent, the less likely he is to challenge it, as this would threaten his own cozy situation. On the other hand, if duopoly and free entry profits are relatively similar, say in an industry with other entry barriers, the competitor is more likely to challenge existing patents, unless strategic considerations outside of our model are at play. The proposition also captures the central finding of Farrell and Shapiro (2008): The more competitive the industry is, i.e. the lower υ_B^C, the less likely it is that a patent will be challenged by competitors and the more pernicious the role that bad patents can play.

Condition (1) gives us a straightforward upper threshold for patent quality, above which B prefers inaction to the challenging subgame, as his expected profits from the latter become negative. This threshold is lowered as κ increases. Thereby, the set of p_B shrinks. For large court costs, the two bounds (1) and (iii) bypass each other and the set of priors for which challenges to patent validity arise becomes empty.[15] The prior regarding patent quality can be interpreted as a measure of patent quality. Therefore we find an alternative mechanism to Chiou (2008) that leads to the result that competitors will contribute to the control of patent quality privately only for intermediate patent quality. In our

[15]The set is given by $\frac{\upsilon_B^D - \upsilon_B^C + \kappa}{\upsilon - \upsilon_A^D - \upsilon_B^C + \kappa} < p_B \leq \frac{\upsilon_B^C - \kappa}{\upsilon_B^C}$.

model, if the patent quality is "too high", there will be no challenges and if the patent quality is "too low". Our model therefore incorporates the central results from Farrell and Shapiro (2008) and Chiou (2008). Let us next consider the analogue to proposition 1 to discern the cases in which the infringement subgame arises:

Proposition 2: *In the Perfect Bayesian Equilibrium of the litigation subgame, B will prefer infringing an existing patent to challenging it and inaction iff one of the following holds:*

(A) Condition (2) is violated.

(B) (a) Condition (2) is satisfied, (b) $\upsilon_B^D - p_B(\upsilon - \upsilon_A^D + \kappa) > 0$*, and (c)* $p_B < \frac{\upsilon_B^D - \upsilon_B^C + \kappa}{\upsilon - \upsilon_A^D - \upsilon_B^C + \kappa}$

Again, the proof follows directly from lemma 1 and 2. (B) is simply the opposite case from proposition 1, where (b) ensures positive expected profits and (c) lets infringement be more attractive than challenging the existing patent. The case (A) appears rather simple, but it signifies one of the possibilities of the patent system "breaking down". If condition (2) is violated, court costs are so high in relation to the monopoly profits that the (certain) payoff of holders of valid patents from going to court is lower than from accommodating the infringer. As a result, the competitor can infringe with impunity. This problem is clearly alleviated with the British system of assigning court fees, in which the losing party has to cover the winner's costs and this would be one important reform suggestion to rejuvenate the American patenting system. But one has to be careful about a too narrow interpretation of κ. For us, this term does not only cover the fees of the courts and lawyers themselves, but in addition also costs of the time for preparation and the hassle of proceedings. Streamlining and simplifying proceedings, as is the goal of the EPLA, would therefore be a method of reducing κ even in the European or British fee system.

These two propositions allow us to consider how the profits of the various player types develop as the average quality of a patent decreases, given its value

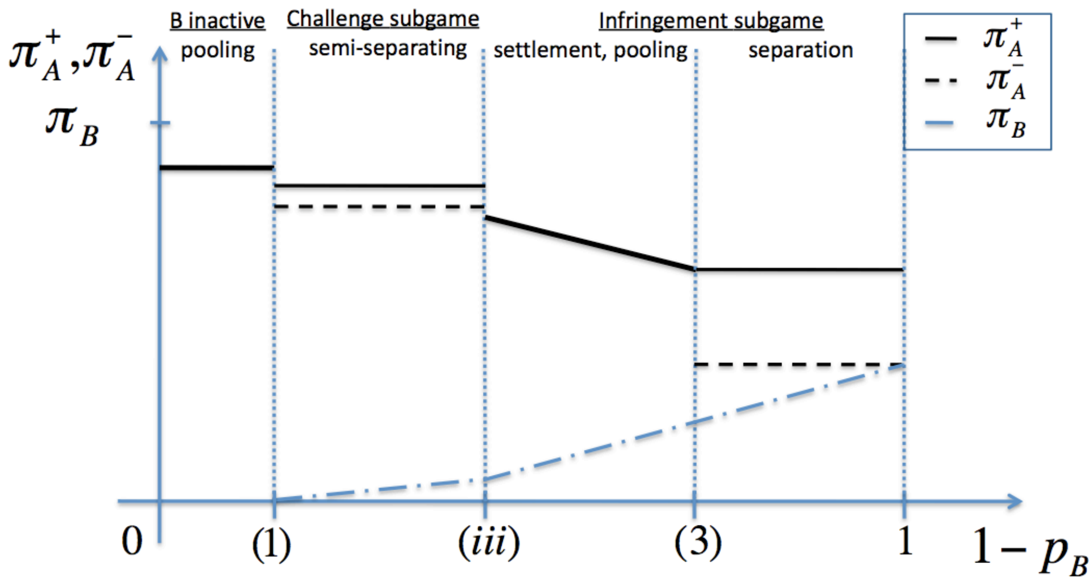

Figure 2.4.: Expected profits of patentees with valid patents (π_A^+), invalid patents (π_A^-) and the competitor (π_B), depending on the prior probability of patent (in)validity.

υ and the corresponding structure of payoffs. **Figure 2.4** illustrates the most complex case, for the case that condition (2) is not violated.[16] As the average quality of the patent decreases (or the share of invalid patents increases) starting at 0, the equilibrium switches from the pure monopoly case to a challenging equilibrium (when condition (1) is satisfied), then to the infringement with settlement (when condition (iii) from proposition 1 is satisfied), to the separating infringement case (when condition (3) is satisfied). As discussed above, the area between (1) and (iii) does not necessarily exist, as infringement can dominate challenging from the view of the competitor.

The figure also allows us to analyze under which circumstances private efforts by competitors may act as complements to the examination process of the PTOs. This will be the case in a socially desirable way if the different forms of litigation act as a deterrent to holders of not-patentable ideas specifically, i.e. if the expected profit of the holder of an invalid patent is strictly lower

[16]If condition (2) is violated, all parties obtain expected profits of $\upsilon_{A,B}^D$.

than the expected profit of the holder of a valid patent. Such a wedge between the expected profits of holders of valid and invalid patents only occurs in the sections to the right of (3), where the patent holders are forced to go to court against infringers, and between (1) and (iii), where competitors weed out a share of the bad patents through challenges.

Overall, the profits of patent holders – both valid and invalid– are weakly decreasing in the share of invalid patents in general.[17] In the following, we will focus on the case $\frac{\partial \pi_A^-}{\partial p_B} \geq 0$, in order to avoid having to deal with multiple equilibria later on. All of the following results can also be derived without this assumption, as long as the profits of both kinds of patentees are weakly decreasing beyond some $\overline{p}_B$ and reach their minimum beyond this point, which is always the case due to our assumptions regarding the specification of industry profits: As the prior probability of patent validity $p_B \to 0$, $(\pi_A^-, \pi_A^+) \to (\upsilon_A^D, \upsilon - \kappa)$, which is the minimum level the respective expected profits can reach. Therefore focussing on monotonously decreasing expected profits greatly simplifies proofs in the following, without affecting the generality of the argument.

Figure 2.5 displays the effects of an increase in the private value υ of a given patent on the profits of the various players. We observe two things: First, profits increase *within* each outcome for all players. Second, the boundaries of the outcomes shift, most notably, (1) is relaxed so that the threshold for challenging a patent decreases. The other clear effect is that (3) shifts downwards, so that fewer cases reach settlement given infringement and it is more likely that infringers will be sued. The effect on the boundary between challenging and infringing patents is unclear given the assumptions made and depends on the signs and relative magnitudes of the changes in $\upsilon_B^D - \upsilon_B^C$ and $\upsilon - \upsilon_A^D - \upsilon_B^C$. We can note as a result, that as the value of a patent increases, patentees are less likely to enjoy the full monopoly benefits therefrom. Further, if patentees suffer losses from an increase in υ, this is either due to the fact that while less

[17]The only exception is if the sum of competitive and duopoly profits more than exceeds the monopoly profits, to be precise, whenever $\upsilon < \upsilon_A^D + \upsilon_B^D + \upsilon_B^C \frac{\upsilon_B^D - \upsilon_B^C + \kappa}{\upsilon - \upsilon_A^D - \upsilon_B^C + \kappa}$. Then, the holder of invalid patents receive a higher expected profits under the pooling and infringement case than under challenges from competitors.

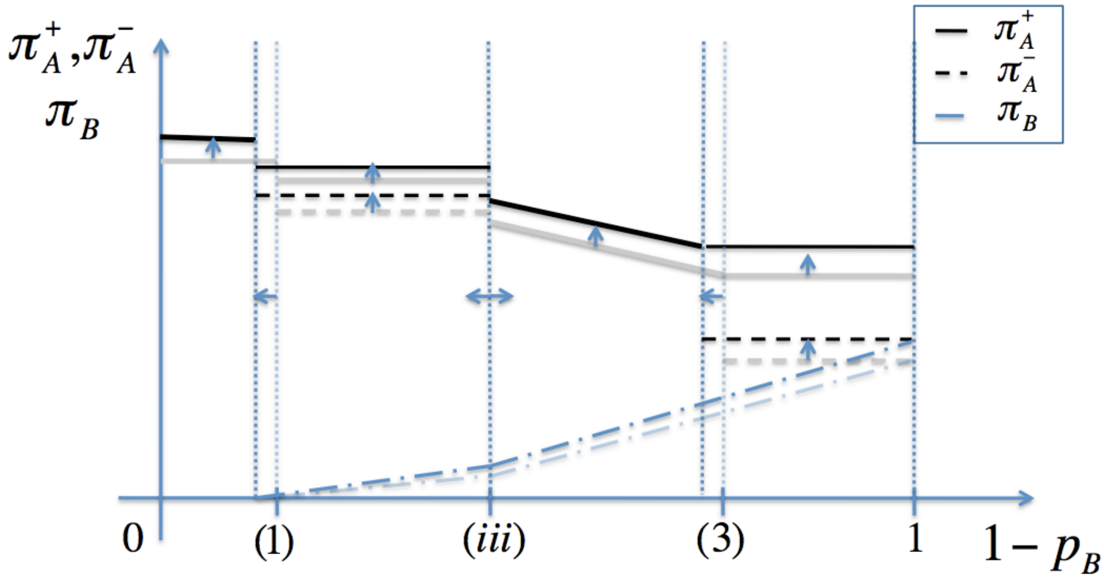

Figure 2.5.: Changes in the payoff structure of the litigation game given an increase in the private value of the patent υ.

profitable patents were not challenged more profitable patents are, or because while less valuable patents were challenged, more valuable ones suffer from infringement. Most importantly, though, as noted above, the expected profits of the holders of valid patents are bounded from below at $\upsilon - \kappa$ (or υ_A^D, if (2) is violated) and the expected profits of holders of invalid patents are bounded from below at υ_A^D. These bounds are strictly increasing in υ.

Incentives for Patenting and Idea Generation

Let $\pi_A(p_B, \iota, \upsilon)$ denote the expected outcome from the litigation subgame for A for a given idea and given prior beliefs regarding the patent quality if granted. Obviously, A will patent this idea only if the following condition holds:

$$\phi(\iota)\pi_A(p_B, \iota, \upsilon) + \delta(\Phi)\upsilon \geq \tau_1 \qquad (2.4)$$

A benefits twofold from a patent application. First, there are the expected profits from obtaining the patent, taking the probability of doing so $\phi(\iota)$ into account if the idea is not objectively patentable. Second, A receives the patent pending status to his idea in any case, which is potentially valuable. If the sum of these two factors exceeds the costs of the patent application τ_1, A will decide to patent. At this stage, there are two potential sources of difference between the expected gains from patenting objectively patentable and non-patentable ideas:

1) $\phi(\iota < \mu) \leq 1$, i.e. non patentable ideas are (weakly) less likely to obtain a patent

2) and $\pi_A^- \leq \pi_A^+$, i.e. given that a patent was awarded, the benefits derived from the patent are weakly smaller for holder of latently invalid than for holders of valid patents. We described above in detail when this second inequality is strict, as only in this case private litigation is complementary to the inspection efforts of the PTO.

Up until now, we have not discussed where the prior probability of patent validity p_B originates. We impose that this prior is correct in the sense that it resembles the actual share of invalid patents in the patent population. Based on this, we are able to state the following Lemma:

Lemma 3: *For a given structure of market payoffs and level of υ, either no idea is patented, every idea is patented, or a cutoff level $\hat{\iota}(\upsilon)$ exists with the following property: For each idea, if $\iota \geq \hat{\iota}$, A applies for a patent and refrains from doing so otherwise.*

Proof: Rearrange (4) to $\phi(\iota)\pi_A(p_B, \iota, \upsilon) \geq \tau_1 - \delta(\Phi)\upsilon$. Note that the right hand side of the inequality is a scalar for given values of υ. Denote the share of valid ideas among all ideas for a given υ as p_{min}. Consider the lefthand side of the inequality. As $\pi_A^+ \geq \pi_A^-$ and it is non-decreasing in p_B, obviously if $\pi_A^+(p_B = 1, \upsilon) < \tau_1 - \delta(\Phi)\upsilon$, no idea will be patented. Analogously, if $\phi(\iota = 0)\pi_A^-(p_{min}, \upsilon) \geq \tau_1 - \delta(\Phi)\upsilon$, every idea will be patented. Our requirement with regard to the priors p_B being correct means that the original prior probability of facing a valid patent $p_B(\Phi, \upsilon)$ is determined by the actual shares in the patenting decision, i.e. the following condition must hold in equilibrium: $p_B(\Phi, \upsilon) = \frac{\int_\mu^1 f(\iota) d\iota}{\int_{\hat{\iota}}^1 \phi(\iota) f(\iota) d\iota} \; \forall \upsilon$. It is simple to show that if B applies for a patent for an idea with ι_1, and there is a second idea with equal value υ and greater quality $\iota_2 > \iota_1$, then B will also apply for a patent for the second idea – the expected value from obtaining a patent is strictly increasing in ι. Analogously, if B derives a negative expected value from patenting an idea with a objective quality ι_1 and there is a second idea with the same value and objective quality $\iota_2 < \iota_1$, i.e. $\phi(\iota_1)\pi_A(p_B, \iota_1, \upsilon) < \tau_1 - \delta(\Phi)\upsilon$ and $\iota_2 < \iota_1$, then $\phi(\iota_2)\pi_A(p_B, \iota_1, \upsilon) < \tau_1 - \delta(\Phi)\upsilon$ and B will also decide not to apply for a patent for the second idea. Combining these statements completes the proof of the Lemma.

Logically the following four cases may arise: **(i)** it is not even profitable to patent valid ideas, **(ii)** it is even profitable to patent the least patentable ideas, **(iii)** it is profitable to patent invalid ideas, as long as their quality is above a certain threshold as well as valid ideas and **(iv)** it is *only* profitable to patent objectively patentable ideas.

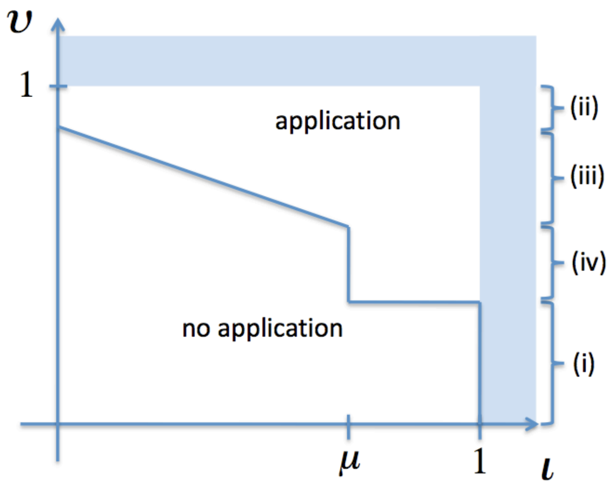

Figure 2.6.: Illustration of Lemma 3 depicting an example in which all 4 cases occur.

Figure 2.6 illustrates Lemma 3 and depicts an example for how patenting decisions are made in (ι, υ) space. It is worth while to consider which of the depicted cases are generic and which only arise under certain circumstances. First note that for every patent application fee $\tau_1 > 0$, case (i) will arise for low enough values υ of valid ideas, i.e. no idea that yields relatively low private benefits will ever be patented. Next let us focus on case (iv), in which only valid ideas are conferred to the patent office. This case will arise if $\lim_{\iota \to \mu} \phi(\iota) < 1$, for then there is a wedge between the value of patentable and non-patentable ideas. Note that if the share of invalid patents is "small enough" $\pi_A^+ = \pi_A^- = \upsilon$, only as the share exceeds the threshold given by (1) a difference between the two values arises. Case (ii) occurs if $\phi(\iota = 0)\pi_A^-(p_B = 0, \upsilon) \geq \tau_1 - \delta(\Phi)\upsilon$, therefore it disappears if τ_1 is large enough. This case is problematic from a policy perspective, as this means that the patent office is unable to deter *any* bad applications. It is more likely to occur for higher private values and larger returns from pending patents. Intuitively, if obtaining a patent for an idea will enable a firm to make an absolute killing in the market, the firm is going to take the comparatively cheap gamble – it only incurs the fixed patent fee – of applying for a patent, no matter how small the chance is of obtaining it. The higher the gains from the patent-pending status, the more relevant this case becomes. (iii) arises in the remainder of constellations. The frontier in (iii) is generally not linear as depicted.[18] Due to the shifts in the boundaries of the litigation outcomes as υ increases as discussed above, the frontier can even be increasing, but only locally.

Finally to close the model, A generates the $(n+1)$th idea as long as the expected profit is larger than its costs, i.e. the following condition holds:

$$\int_0^1 \int_{\hat{\iota}(\upsilon)}^1 [\phi(\iota)\pi_A(\upsilon, \iota) + \delta(\Phi)\upsilon - \tau_1] f(\iota) g(\upsilon)\, d\iota\, d\upsilon \geq C(n+1) - C(n) \tag{2.5}$$

Therefore also patents that should formally and legally not have been granted

[18] Instead this resembles a case in which $F(\iota)$ is independent of υ and the expected profits from receiving a patent are either constant or a linear function of υ. This occurs for example if the same litigation-equilibrium arises independent of υ.

give inventors an incentive to exert innovation efforts as long as they generate positive expected payoffs. For example, a laxer inspection policy will result in the generation of more ideas. This is an obvious simplification, as it does not take second order effects such as patent-thickets etc. into account. Nevertheless, we have created a relatively complex system. For this general setting with general distribution functions, a closed-form solution does not exist. But despite this complexity, in the following section we will be able to generate a number of important insights concerning the proposed piecemeal reforms and adaptations of existing patent regimes, focussing on pointing out trade-offs and problems that have been mainly overlooked so far.

2.4.2. Comparative Statics - Potential Approaches to Reform

Application Fees

One obvious difference between the European and the US patent regime is the significant difference in fees that are imposed for patent application and maintenance. Various authors, most recently Bessen and Meurer (2008a) and Chiou (2008) have suggested to increase patenting fees in order to deter lower quality inventors and raise the average quality of patent applications.

Our setup allows us to dissect the results of such a reform in an extremely straightforward and simple manner. Let us consider a change from τ_1 to τ_1'. This affects the patenting decision at two margins. As shown above, the lefthand side of the condition that implicitly defines $\hat{\iota}$, i.e. $\phi(\iota)\pi_A(p_B, \iota, \upsilon) + \delta(\Phi)\upsilon \geq \tau_1$, is strictly increasing in ι given that $\iota < \mu$ for each value of υ. Therefore if one adds a constant to the righthand side, the cutoff value $\hat{\iota}(\upsilon)$ must increase as well, which resembles a shift of the frontier in Figure 6 to the right. We will denote the new cutoff value as $\hat{\iota}(\upsilon)'$. Intuitively, this is the desired effect, as fewer invalid ideas are submitted to the PTO office for scrutiny, and captures the intention of those proposing this kind of reform.

There is a very simple second effect though, one that directly affects patentable ideas only. Consider the case in which only valid ideas are patented, therefore $p_B = 1$ and $\pi_A^+ = \upsilon$. Here the patenting decision is defined by $(1 + \delta(\phi))\upsilon \geq \tau_1$, which directly implies the minimal private value of an idea that is patented is $\underline{\upsilon} = \frac{\tau_1}{1+\delta(\Phi)}$. An increase in the application fee τ_1 moves this frontier upward, which is the effect at the second margin. This second effect is extremely undesirable, as it precludes additional objectively patentable ideas from being patented, because the costs outweigh the *private* benefits.

Further, both effects lead to a strict decrease in the expected profits to be derived from generating new ideas and should therefore lead to less generated ideas (and patent applications) in general. As a short sidenote: Due to the static character of our model, we are overlooking a further potential effect: The smaller number of applications may lead to a decrease in the delay between application and patenting and thereby reduce the benefits from pending patents, which would be a further positive effect. Arguably, this should not happen in the short run, though, due to the "stockpile" of applications that has accumulated over the years and still has to be dealt with. We summarize these considerations in the following proposition:

Proposition 3: *An increase in the patent application fee from τ_1 to τ_1' leads to an increase in the relative share of valid applications only if the following holds:*

$$\frac{\int_{\underline{\upsilon}'}^{1}\int_{\mu}^{1} f(\iota)g(\upsilon)\,d\iota\,d\upsilon}{\int_{\underline{\upsilon}'}^{1}\int_{\hat{\iota}'}^{1} f(\iota)g(\upsilon)\,d\iota\,d\upsilon} - \frac{\int_{\underline{\upsilon}}^{1}\int_{\mu}^{1} f(\iota)g(\upsilon)\,d\iota\,d\upsilon}{\int_{\underline{\upsilon}}^{1}\int_{\hat{\iota}}^{1} f(\iota)g(\upsilon)\,d\iota\,d\upsilon} > 0 \tag{2.6}$$

Proof: The proposition follows directly from lemma 3 and the definition of $\underline{\upsilon}$.

When the condition in the proposition does not hold, the share of valid patents actually *decreases* with a change of the application fee! As the condition is not necessarily intuitive in this form, we may benefit from considering the marginal effects only. From our previous discussion it is clear that a change in

τ_1 will affect the mass of valid as well as the total mass of applications. Let us denote the original share of valid patent application as $\gamma(\tau_1)$ and the total mass of applications as $\alpha(\tau_1)$. Then simple calculus allows us to derive the following much more interpretable corollary to proposition 3:

Corollary to Proposition 3: *The marginal effect of a change of the application fee τ_1 to τ_1' on the relative share of valid applications will be positive only if the following holds:*

$$|\alpha'| > \frac{(1 - F(\mu))g(\underline{\upsilon})}{1 + \delta(\Phi)}(\gamma(\tau_1))^{-1} \tag{2.7}$$

From this one can easily determine that the average quality of patent applications is more likely to decrease for increases in the application fee, if the density of patents with the marginal private value $\underline{\upsilon}$ is relatively high. This stems from the fact that on the lower margin only valid ideas are patented. This effect mechanically becomes stronger, the lower the original patent quality is, as the decrease in valid patents then has a stronger weight, as well as for laxer patent regimes (lower levels of μ). As a direct result of this, in regimes with a relatively high average quality of patent applications such as the European patent office, a *decrease* of the patenting fee could increase the average quality of applications. On the other hand, for the US, in which the complaints about the quality of applications are widespread, an increase in the patenting fee should be less problematic.

Intuitively, there is a kind of "positive selection" at work here. As the private value of patents decreases, the quality cutoff level for which patents are profitable increases, until at the margin only the best patents survive - if an inventor expects a patent to be of relatively low private value to him, he must be relatively (or completely) sure of its patentability to be willing to incur the costs of patenting. Raising patent fees therefore generically kills off high-quality patent application over-proportionally.

Litigation Costs

Various developments in the US and Europe have affected (or will affect) the way patent litigation functions. In the US, the evolution of a specialized appellate court has reduced legal uncertainty (which can be interpreted as a form of cost in our simplified model) and possibly led to or at least encouraged an increase in patent litigation.[19] Concurrently, in Europe the European Patent Litigation Agreement has been seeking to establish a unified European Patent Court with the expressed goal of reducing patent litigation costs.[20] Our model setup allows us to dissect the effects of such a decrease in some detail.

Let us first consider the changes in the litigation subgame. We collect our qualitative results in **Figure 2.7**. Again we observe two related kinds of shifts: The first affects the the payoffs of the individual players directly, and the second, as a result, moves the boundaries between the different litigation scenarios. In the infringement game, we get the (well established) result, that the likelihood of settlements will *decrease* if the litigation cost are reduced. From a competition policy and welfare standpoint, this is a highly desirable result, as it reduces the scope for collusion. In our figure, this is represented by a shift of boundary (3) to the left.

Next, the reduction makes challenging existing patents more attractive to the competitor. This effect is unambiguous at both margins. For relative high priors of patent quality, challenges occur where they did not before. This is represented by a shift of the boundary at (1) to the left. But also at the margin between challenging and infringement/settlement, the effect is clear, with challenging becoming relatively more attractive compared to infringement. Therefore the boundary at (iii) moves to the right. Each of these effects is arguably desirable from a social welfare point of view.

The payoffs of the various player-types develop in an interesting pattern. The competitor is weakly better of as court fees decrease – the threat of litigation

[19] See Cook (2007) for more details. Further see Bessen and Meurer (2008a) for a very critical assessment of the recent developments of the patent judiciary in the US.

[20] See for example http://www.epo.org/patents/law/legislative-initiatives/epla.html for details.

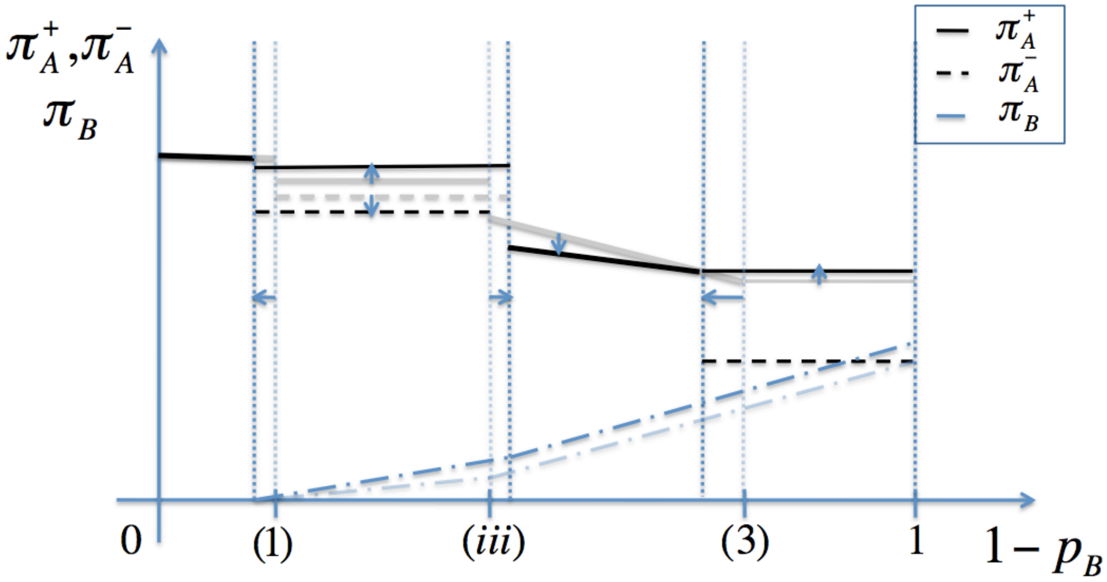

Figure 2.7.: Changes in the payoff structure of the litigation game given a decrease in the costs of going to court κ

bears slightly less weight in the infringement case (which also makes settlement less likely) and the payoff of a challenge, whether successful or unsuccessful is strictly higher. The change to the expected payoff of a holder of a valid patent is somewhat ambiguous. Given that his patent is challenged, the lower court fees make him better off, yet he will be challenged in more cases. It is this effect that Bessen and Meurer (2008a) worry about in the face of litigation reform in the US. Patent litigation becomes more likely also in the infringement subcase, as argued above - as settlement is not necessarily desirable from an antitrust perspective, as has been brought to attention by Shapiro (2003), this could be a beneficial effect, but it decreases the expected payoff of patent holders. Finally, in the cases in which the holders of valid patents have to sue to protect their investment, i.e. to the right of (3), the decrease in costs strictly increases their profits.

The effects on the profits of holders of invalid patents are less ambiguous. The fact that more patents will be challenged strictly decreases the area in which pooling equilibria arise: The shift to the left of (1) implies that more patents will be subject to challenges, from which bad patents suffer more strongly than good ones. Further, the shift of (iii) to the right implies that patents that would otherwise have been infringed and settled are now subject to challenges as well. While the effect on the profits of holders of bad patents in this area might even be positive, as discussed above, nevertheless this increases the area in which a wedge exists between good and bad patents and thereby increases the likelihood that the courts may be able to aid the work of the PTO. Finally, the shift of the boundary (3) to the left implies that holders of good patents must defend more claims in court that they would otherwise have settled. As the holders of bad patents never go to court, this reduces their income from settlements and again increases the wedge between the payoffs to the two types.

To summarize, this analysis shows that streamlining court-proceedings can be poison or medicine, depending on what you believe ails the patent system. If, like Bessen and Meurer (2008a) one believes that the main problem of the system is the erosion of payoffs to patenting from too much litigation, then the effect of lowering costs is unclear – not strictly negative, as one might have expected. On the one hand, it should lead to more litigation, but on the other hand the costs to patentees decrease, so the net effect is unclear. At the same time, the reform steps in this direction unambiguously strengthen the role that litigation and courts can play in sifting out bad patents from the system and might therefore reduce the burden on the PTO.

Extension: Diligence of the Patent Offices

The effects of an increase in the diligence of the PTO appear obvious at first sight. An increase in the inspection intensity should lead to more bad applications being rejected, so that it becomes less attractive to try to pass of bad ideas as patentable. Consider an adjustment of the patent office's policy such that Φ changes to Φ' and thereby $\phi(\iota)' < \phi(\iota)\ \forall \iota < \mu$. By itself, this effect would be purely beneficial. But as assumed above, this improvement comes at a cost, as the returns from pending patents increase to $\delta(\Phi') > \delta(\Phi)$. Note that an increase of the required diligence in the PTO's assessments would in reality not only affect new patent applications but also the existing backlog of past applications that has accumulated.

Revisiting condition 2.4, we had established the following regarding the decision whether or not to apply for a patent with a given idea:

$$\phi(\iota)'\pi_A(p_B(\Phi'), \iota, \upsilon) + \delta(\Phi')\upsilon \geq \tau_1 \tag{2.8}$$

The effects of this reform then are far less clear cut than we would have expected. Note first that in any case the cutoff level $\underline{\upsilon}$ decreases, as the return from valid patents is increasing in δ, which may be considered a positive side-effect. The overall outcome on the other margin depend on the relative sizes in the changes in patenting probability and profits received during the pending phase. Instead of over-stretching our model at this point with regards to the assumptions we have made (for example the linearity of pending-revenues in the private value of an idea), we would like to point out that the profits during the pending phase have a fundamental effect on the current patenting system, yet there are next to no empirical findings with regard to this topic, while they are clearly sorely needed.

2.5. Discussion and Outlook

The current regime of issuing and enforcing patents and the resulting rights of patent owners has faced increasingly harsh criticism in the recent past. There have been calls for completely abolishing the system or at least carrying out a fundamental reform thereof. Due to the enormous investments of the holders of existing patents, it appears relatively unlikely that such a fundamental reform is going to happen soon.

What appears far more likely, instead, is a continuation of the current praxis of minor piecemeal improvements (or rather changes, to phrase it more neutrally). We add to the understanding of the effects of three of the most considered – or currently being implemented – steps: An adjustment of the application fees, simplified and cheaper adjudication of patent conflicts and finally a policy of stricter scrutiny by the PTOs. Our approach is new in two central regards - first, we take three major aspects of the current system into account in our theoretical model, i.e. policy setting, the role of the PTO, and the interplay of patentees and potential challengers/infringers in the market and in the courts. Second, we take into account that ideas come in a plethora of form and shape and that they may vary continuously in their private value to the patentee and in their patentability (which to a certain extent reflects their value to the community).

From our approach, we gained both new insights and new questions. Allowing the competitor to both challenge or infringe upon an existing patent, we found that below a certain threshold of patent quality, the competitor loses her interest in policing patent quality through challenges but will prefer to infringe upon the rights of the patentee instead. This effect is strongly exacerbated in industries in which there are relatively few barriers to entry - a patent does not only protect its holder. In this context, there is clearly a thin line between tolerating an infringer and anticompetitive settlements of weak patents - focusing on explicit agreements from an antitrust perspective does not do justice to this issue.

Allowing for continuous types allowed us to derive a surprising result with regard to an often contemplated reform step: Raising application fees may actually decrease the average quality of patent applications. This will especially be the case if the distribution of private patent values is skewed towards lower values - which appears to reflect reality very well - and if the standards for patenting are relatively lax. Further, we showed that a decrease in litigation costs through specialized courts (or a centralized court in the case of the EU) may actually make the work of the patent office easier, by forcing a wedge between the payoffs from good and bad patents (or increasing an existing wedge). Further it can destabilize the infringement "agreements" discussed above.

Finally, we introduced a problem into our setting, which is so far very little understood - patents yield benefits to their owners before they are approved by the PTO. Prior to the decision, even applications that are later rejected may be of considerable monetary or strategic value to their owners. We strongly appeal to empirical economists working in this field to help better understand the issues connected to the “patent pending” stamp. The second empirical issue raised is the interplay between market structure (e.g. barriers to entry) and the outcomes of patent disputes. Our model predicts that the openness of an industry should be negatively related to the probability of a patent being challenged.

The approach of this article was to make relatively few assumptions and see which general results could be derived therefrom - clearly, one logical next step would be to breathe more life into the market setting by imposing tested forms of competition that are well understood in the spirit of Farrell and Shapiro (2008) together with specific distributions in order to be able to work with closed form solutions. We believe that this more flexible setting is a valuable complement to the existing literature, as it enabled us to point out a number of tradeoffs that are worth studying in more detail.

3. Individual (ir)rationality? Behavior in an emerging social online-network

3.1. Introduction

3.1.1. Motivation

Actions and interactions in social networks involve highly complex decisions. The resulting payoffs can be immense, such as acquiring a job offer that one would otherwise not have received or founding long-lasting mutually beneficial partnerships. Regarding the timing, these benefits or detriments can be immediate or gradual, lasting or short-lived. Their source may be linked to an individual directly or through any number of intermediaries. While the importance of networks in highly relevant areas, such as the spread of medical knowledge (see Coleman et al. (1957)) or the likelihood of obtaining a desirable occupation (see Granovetter (1974)), has attracted the attention of social scientists many decades ago, researchers for a long time grappled with the empirical and theoretical problems raised by their complexity.[1]

The last decade has seen major breakthroughs in the way economic theorists think about and approach the subject of networks. Works such as Jackson and Wolinsky (1996), Bala and Goyal (2000) and Jackson and Watts (2002) have made many of these concepts tractable in a graph-theoretic setup. In these articles, rational agents maximize the payoffs generated from links to others, which in turn allows predictions concerning the structure of the underlying networks. Yet in many social networks one encounters potentially emotionally

[1]See Granovetter (1976) for an early discussion on the empirical issues raised by the dimensionality of data.

charged terms, such as "friendships" in social online-networks or "partnerships" in business relations. Consequentially, a more recent strand of the literature has realized that concepts such as inequality aversion or trust may play an important role in interactions between agents in social networks. Experimental economists have therefore introduced these concepts into simple network models and tested their behavioral predictions under highly controlled settings, providing evidence that often strongly contradicts the "rational" predictions for equilibrium network structure.[2]

The contrast to these controlled settings highlights some of the challenges that non-experimental empirical work on individual behavior in social networks faces: e.g., it is impossible to disentangle the motives of individual persons or entities entering into relationships, costs and utilities of decisions are mostly unobservable and may differ substantially between individuals, etc. Most importantly, though: There is next to no data that combines micro- (individual decisions and characteristics) and macro-level (structure of the network overall, locations of links inside a network) observations.

Our study aims exactly at this gap, using a completely unique data-set. We were able to obtain data from an emerging online-social network: on the one hand on user behavior of more than 30,000 registered users, for example with regard to provision of a public good type effort, number of private messages sent, number of public comments posted, entry into new and severance of existing "friendships", and on the other hand on the structure of the underlying network itself, i.e. number and location of friendships between users and overall network size.[3] We use this data first of all to perform what can be considered "basic" research: Addressing the challenges by experimental studies, we demonstrate clearly that users behave in ways that are compatible with the rational paradigm. Based on these findings, we address two further issues that could

[2] See Kosfeld (2003) for a survey of the early literature in this vein.

[3] As part of the agreement concerning the data, the name of the company and its service is not used in this study. All user-related information was provided to us in a way to completely assure users' confidentiality and protection of their personal data. We were only provided with randomized user-IDs that cannot be linked to individuals in any way and we have no access whatsoever to individuals' personal characteristics.

barely be studied empirically so far:

The first is the provision of local public goods by users. Users in the network can provide and organize information – as described below in detail – two kinds of activity from which both they themselves and their direct friends derive utility. Standard theory, such as Bramoullé and Kranton (2007) in the specific context of a network with information provision, would lead us to expect to find evidence of free riding in the sense that users should provide less effort themselves if they have friends who provide more effort, all else given. Interestingly, we find only weak evidence for this type of free riding. Instead – and far more pronouncedly – users seem to react in a positively reciprocal way to their friends' effort provision - that is they react to their friends' provision of effort by exerting more of their own, a phenomenon that has been noted before in the context of peer-to-peer networks among others by Gu et al. (2009).[4]

This leads us to the second issue of our study: Direct, negatively reciprocal behavior as an enforcement device. Studying the formation and stability of individual links, we find strong evidence that users are willing to incur costs to actively punish free riding by their friends by severing (otherwise purely beneficial) relationships.

Our findings are highly relevant, as the provision of public goods in Web 2.0 services is creating tremendous, constantly increasing values. Microsoft's partial acquisition of Facebook priced the latter company as a whole at multiple billion dollars[5] - which is somewhat surprising, because apart from the infrastructure, the entire content of the social-networking site as well as most applets are created and provided by users. The same holds true for 3 other purely social networks that together with Facebook form 4 of the 20 most frequented sites on the internet.[6] If one further considers massively popular sites with a strong social-network character such as Youtube.com or flickr, it becomes obvious that social online-networks are encroaching on many people's daily lives and that a major building block of this development is free user-provided content. In

[4]For an introduction to the economics of reciprocity, we refer to Fehr and Gächter (2000).

[5]See, for example http://www.time.com/time/business/article/0,8599,1675658,00.html

[6]The sites are MySpace, FaceBook, Orkut and Hi5.com, according to alexa.com as of April 2010.

the course of this article, we seek to provide new economic insights into the behavior - and potential motivation and payoffs - of individuals in these kinds of networks.

We proceed as follows: In the following subsection we present the most closely related literature. First the theoretical foundations of the theory of social networks, then the "experimental attack" and finally we briefly cover existing empirical studies. In **Section 2** we introduce and describe the online-network we wish to study as well as the features that make it especially interesting from an economic standpoint. **Section 3** relates these features to some simple theoretical concepts that will form the basis for our further inquiry. **Section 4** contains the data description and our empirical analysis. **Section 5** concludes and proposes resulting new venues for future research.

3.1.2. Related Literature and Contribution

The literature relevant to our questions can be separated into theoretical, experimental and empirical studies.

From the **theoretical** perspective, a set of by now canonical articles has enabled economists to come to term with the complexity of the subject, allowing them to make predictions about the structure of networks and to judge their efficiency. Perhaps the initiators of this development are Jackson and Wolinsky (1996). Their graph theoretic framework, depicting individuals in a network as the nodes in a graph and the connections or links between these individuals as the arcs, has become the de-facto standard approach. Individuals derive utility from the links they are involved in (and indirectly also from the links of those they are connected to), while it is costly to maintain direct links. A network is "stable" if the utility of agents cannot be increased by creating a new or severing an existing link. The predictions concerning stable networks from this model range from complete interconnection (high benefits to costs ratio) over a "star network", in which one node is connected to everyone else, while the peripheral nodes maintain only a single connection each, to a completely disconnected network (lowest benefits to costs ratio). Bala and Goyal (2000) enrich the setup by explicitly modeling the linking-strategies of players in a non-cooperative game and distinguishing between cases in which information (and thereby benefits) flow in only one or both directions. In addition to the stable network structures mentioned above, they find that a symmetric circle or wheel network in which each agent is linked to exactly two other players. Of further importance in their setup is the concept of decay of benefits - the more intermediate links separate two indirectly connected players, the lower the benefits they derive from each others' information.[7]

There are a number of important extensions to these basic models. Goyal (2005) notes that depicting the existence of a link between research affiliates, for example, as a purely binary variable may often be too much of a simplification. Instead, partners can invest specifically into the stability of individual links, so that heterogenous link strengths may arise in equilibrium. Bloch and Dutta (2009) adopt this concept and model it explicitly. They find that under relatively broad conditions, the stable and efficient network again will have a star structure.

Three papers of which we are aware of try to model the provision of public goods in social networks. In Cho (2006) ex-ante symmetric agents agree to a binding contract that covers both the amount of a local public good each agent provides as well as the form of the network through which the benefits of the public good are transmitted. In this setup, equilibria exist, in which agents almost surely provide efficient levels of effort if they are patient enough. In Bramoullé and Kranton (2007) agents only specify their effort levels, taking the structure of the network as given. Effort provided by an agent's neighbors is a substitute for her own efforts. They show that under these assumptions free riding exists in every social network in equilibrium and adding links between players reduces individual incentives to contribute. In the model of O'Dea (2008), depending on the outcome, one can find both equilibria in which agents

[7]For an excellent survey of the theoretical literature on network development and stability, see Jackson (2004).

voluntarily provide public goods and equilibria which resemble cost sharing for an excludable club good.

These clear and distinct predictions concerning network formation and structure have piqued the interest of a number of **experimental** economists, who find significant deviations from the rational depictions. Falk and Kosfeld (2003) directly test the hypotheses developed in the Bala and Goyal (2000) model and show that in settings in which the Nash-equilibrium provides for symmetric payoffs and little coordination efforts, the predictions are relatively accurate. Yet in settings that should lead to star networks, with relatively high costs per link and high information benefits, the rational equilibrium predictions have almost no explanatory power. They show that this is only to a small extent explained by the higher requirements with respect to coordination. Instead they argue that star networks provide users with extremely unequal payoffs and therefore the deviation from the rational equilibrium may be explained by subjects' inequality aversion. Cagno and Sciubba (2007) attempt to dissect the various drivers of individual decision in network formation. They identify two conflicting driving forces: Optimizing best-response behavior only explains the observed behavior to a certain degree. Instead they find that reciprocity and inertia play important explanatory roles. In a highly surprising study, Cagno and Sciubba (2008) combine a network formation experiment with a trust experiment. They find that the levels of trust in a trust game played after an endogenous network formation game are significantly *lower* than in the simple trust game. Subjects hereby tend to trust players less on average with whom they have interacted previously than they trust complete strangers. The authors deduce from this that potentially existing reciprocal behavior in the network formation game does not foster trust. This indicates that there are psychological mechanisms at play in network interaction that still have to be explored.

Not all experimental studies have been purely detrimental to the rational behavior paradigm. Corbae and Duffy (2008) test the rational behavior hypothesis in rather specific experimental setting, in which individuals can form risk-sharing partnerships through a non-cooperative proposal game with idiosyncratic risk. They find that individual behavior is generally well predicted by rational theory, which agent achieving near-efficient outcomes. Goeree et al. (2008) analyze a setting in which agents with heterogeneous linking costs can form links under different informational settings. While for homogenous agents equilibrium predictions cannot be validated, the heterogeneous agent predictions are well suited to explain individual linking behavior. Interestingly, star networks do develop in this experimental setting as frequently as expected.

There is a large number of **empirical** studies regarding networks in general that mainly focus on locating and quantifying network effects, such as prominently Saloner and Shepard (1995) or Gowrisankaran and Stavins (2004). For a survey of this strain of the literature, we refer the reader to Birke (2008). There is only a limited amount of empirical economic research on social networks, though. Prompted by Newman (2001) who looks at the properties of research collaborations from a physicist's perspective, Goyal et al. (2006) and Rosenblat and Mobius (2004) both analyze the structure of research networks in economics. Goyal et al. (2006) find a collaboration structure composed of a number of stars (authors that write a relatively large number of papers together with coauthors) and a large periphery (composed of authors who only published papers together with one of the star-authors), which suits the structural predictions of theoretical models. Based on a model that they develop, Rosenblat and Mobius (2004) look at the rise of the internet and electronic communications as a kind of natural experiment which lowered the costs of communication between potential collaborators. They found that the degrees of separation between individuals on average decrease and each author is more likely to coauthor a paper with a distant author from a similar field. On the other hand, the likelihood of coauthoring a paper with a coauthor from a dissimilar field significantly decreases. Further, Kossinets and Watts (2006) use the email-communications inside a large US college to draw up a social network structure and discover various characteristics that determine the likelihood that two individuals within the college will be linked directly.

A small wave of more recent yet to be published papers proceeds in a direction that is almost a hybrid between economic and information systems research - in their subject, these are most closely related to our paper, even though they strongly differ in methodology and focus.

Kumar et al. (2006) try to characterize user types in two huge social online-networks, Flickr and Yahoo!360. They characterize 3 kinds of users: passive members, inviters who entice others to join, and "linkers" who make full use of the social networking capabilities of the services. Gu et al. (2007) look at at peer-to-peer IRC music sharing service and show that it resembles a two-sided market in many ways, in which contribution and consumption are highly complementary.

Three studies specifically scrutinize public good problems and free riding: Asvanund et al. (2004) try to quantify the extent of free riding in a peer-to-peer file sharing network (Gnutella) on a macroscopic level. They find that a substantial share of users (42%) behave as free riders, i.e. access others' files without providing any of their own. Xia et al. (2007) discover that users are more likely to share files in online sharing communities if he has himself benefited from the community or has a recognized social status. Finally Gu et al. (2009) consider the effects of indirect reciprocity on the public good provision behavior of members of a peer-to-peer file sharing network. They find that the propensity of an individual user to free-ride depends on the prevalent behavior of other users in the community, which they construe as evidence for an indirect reciprocal behavior. Further they find evidence for reciprocity as a social norm, as free riders are discriminated against through voluntarily enforced settings on communal servers. These studies are not comparable to ours, as they do not have the individual characteristics of users at their disposal. This allows us to apply individual utility maximization behavior and analyze user behavior in a microeconomic context. Apart from this fact, pathological incentives may arise through the fact that file sharing is illegal behavior, while the service that we study operates well within the realms of copyright law.

Against this extensive background, our contribution is the following: We make use of a unique data-set in which we observe both individual characteristics and behavior as well as the link structure among all individuals within an actually existing social online-network. In contrast to the findings of a series of experimental studies, we provide evidence that network formation and linking decisions are motivated by strategic behavior that can be explained within the utility maximizing rational agent framework.

Then we focus on public (or more precisely club-) good provision within the network. Looking at the entire behavior with respect to the forming and continuation of friendships among users over approximately four months, we show that negative reciprocity, i.e. unilaterally severing existing friendships, even though they are purely beneficial, is used to punish free-riding behavior. Finally, we look at the overall effects of this. Using a user-fixed effects panel estimation procedure, we show that there is only very weak evidence for the existence of free riding behavior, despite an almost textbook public good problem. On the contrary, we are able to identify signs of positively reciprocal behavior. The effect of the latter more than compensates the effect of free-riding that we observe.

3.2. The Social Online-Network under Scrutiny

3.2.1. Function and Features of the Network

In this article, we study an emerging European social online-network, which we will refer to from now on as the Network. Various features make the Network especially interesting for economic analysis. The basic infrastructure is very similar to well-known sites such as Facebook. Each user has an individual profile page which is accessible to others and which allows him to provide others with information about himself, such as age, gender, hobbies etc. As a further feature familiar from other networks, individuals can become “friends” with other users - for this to occur, one of the users has to propose the formation of a bilateral link to another and the second user has to accept this proposal. For legal reasons, the number of friendships a user is allowed to enter into is capped at 150, but this limit is non-binding for all users for the duration of our study. While it takes the consent of both users to become friends (link-formation is bilateral), either of the two can end the friendship at any later point of time (unilateral link severance).

The main differentiating feature or unique selling proposition of the Network is that users can upload digital copies of music that they legally own to an online music library. We call this uploading activity **provision of information**. Via a player embedded in their web-browser, users then can listen to (but not download) music in their library. Further, and this is the main motivation for our study, users can also listen to music in the library of their direct, i.e. first-degree, friends, as long as the friendship lasts. A user therefore has access to all the music that either he or one of his direct friends has uploaded while he is online.

A feature of the Network that is something of a nuisance to users is a blessing to our study: The access to individual music files is somewhat complicated. For one, the music available to a given user is distributed into different music libraries (e.g. music uploaded by himself is one library, but music uploaded by each friend is stored in an individual separate library). The music files in each of these libraries are sorted alphabetically, which makes it potentially time-consuming to find specific songs in large libraries or to find something interesting while browsing through them. Tests with medium sized libraries of around 1,000 songs size showed that it took more than 6 minutes to locate a song beginning with the letter "m". While the service does provide users with a search function, searches due to the programming of the website always cover the library of the *entire* Network, i.e. both files that a user can access and files that he cannot. From the search results (itself a list of hundreds of songs, depending on the popularity), a user can only establish whether she has access to a given file by adding it to her player and waiting for an error message to appear upon hitting the play-button and waiting for the song to load. Due to the extreme usability issues with this procedure, the site was recently completely refurbished, but this took place outside the time-window that we observe.

It is exactly this nuisance that introduces the second kind of effort provision into the equation. There is a mechanism that reduces the inconvenience for users. They can create "play-lists" by adding multiple songs to the embedded music player. Any number of songs can be added from the libraries that a user has full access to. These play-lists are temporary in principle, i.e. they are deleted whenever the browser cache is cleared. But a user can also save and give a name to a play-list that he has compiled. This has two effects: It stores the list with the Network and makes it available to the user whenever he logs in again in the future, as well as to anyone visiting his website. Clearly, this will be most useful to the user's friends, who can access many of the same songs as her due to the greater overlap of their first-degree libraries. For this reason, we call the creation and storing of play-lists **organization of information**. The second effect is that other users, who are not friends yet, cannot view a given user's songs, but they can view his playlists. A user's playlists therefore are the main indication of his value as a source of songs to others in the network.

A user can prevent one of his friends from accessing music-files (only) by ending their friendship. As this also significantly decreases the value of his playlists, provision and (indirectly) organization of information have the characteristics of classical club goods.

3.2.2. Network Size and Development

The web-site came online in December 2007 with 2,000 registered users. This number grew to 9,000 in March 2008, more than 20,000 in September 2008 and exceeded 30,000 in November 2008, when the data for this study were mainly collected. For many - maybe even most - questions of interest concerning online networks, the number of registered users is not a very good indicator, as it may include individuals who only registered at some point of time but never actually used the service or have ceased to do so at some point in the past; therefore we will briefly discuss the most widely used alternative measures.

According to the common practice of sites such as facebook, myspace or similar sites such as the German studi-vz, a user is considered active if she has logged into her account in the course of the past 30 days. At first glance, it is not obvious whether and how this measure is meaningful. **Table 3.1** below sheds some light on this question. In it, we have gathered descriptive statistics for various such cutoff levels of user activity as of October 15th 2008.[8]

	All Users	Act30	Act20	Act10
Friends	0.92	2.16	2.50	3.91
Songs Listened	35.57	100.82	121.50	208.63
Songs uploaded	28.62	50.23	56.93	96.88
Play-lists created	1.63	2.55	2.80	3.80
Messages sent	0.38	1.30	1.67	3.02
Comments left	0.03	0.12	0.16	0.28
Days logged in	2.92	6.03	7.01	10.57
#Observations:	31,455	6,013	4,466	2,291

Table 3.1.: Descriptive statistics for various activity measures

Some important observations can be gleaned from these numbers, both about the activity measures in general and about the Network in particular. First of all, each more restrictive sub-sample includes users that are on average highly significantly more active than the users in the more inclusive sample. For example, users that have been active in the last 30 days (Act30) upload almost twice as many songs on average as the general population, i.e. those users that are left out of the sample are significantly less active. One loses a large number of observations in the process, from all registered users to Act30, the drop is precipitous, from 31,455 to 6,013. These findings indicate that the frequently employed measure does have some merit. **Figure 3.1** depicts the development

[8]Act30 designates user who have logged into their accounts in the last 30 days, the other columns have analogous interpretations. By line, the sample means of the following are reported: number of friends per user, number of songs listened to per user, number of songs uploaded per user, number of play-lists created per user, number of messages sent per user, number of public comments left per user, number of days on which the user has logged into her account at least once per user. All respective sample means differ at 1% significance level (t-tests).

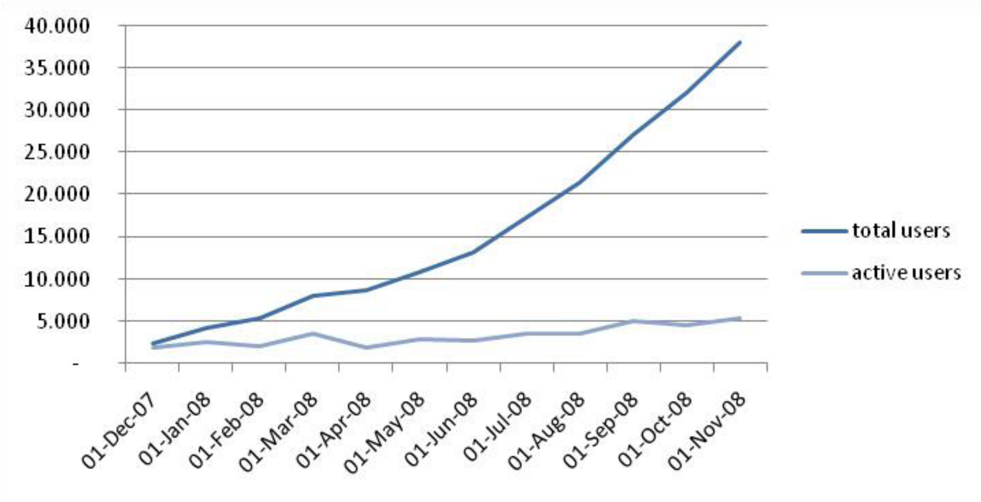

Figure 3.1.: Development of Act30 Users over time.

of registered vs. active users in the period from December 2007 to November 2008. While the number of registered users increases more than 17-fold, the number of active users increases only by a factor of 2.5 in the course of this period.

Two further findings from Table 3.1 concerning the Network are particularly important for our study. On the one hand, note that "social" activities, such as posting comments or sending messages to other users appear to play a very minor role in the network, with even the most active group of users having sent no more than a total of 3.02 messages and posting no more than 0.28 comments on average. On the other hand, listening to and uploading songs take up a much larger share of users' activities, which points in the direction that these music-consumption related activities are the main source of utility for those involved in the Network.

3.3. Analytical Framework

One of the aims of this article is to provide evidence for strategic and utility maximizing behavior among users in the Network. In order to facilitate the understanding of readers and help us formulate our hypotheses more precisely, we will sketch a simple formal model of link formation and individual behavior within the Network. This will form the basic framework for our empirical analysis. We hereby draw on the notation and concepts introduced by Bala and Goyal (2000), Jackson and Wolinsky (1996) and Bramoullé and Kranton (2007). As we do not intend to explain or analyze the equilibrium network structure, but instead individuals' decision behavior given the contemporaneous structure, we will keep the complexity of the exposition to the absolute minimum.

Following the established graph-theoretic approach, we think of the Network as a set of nodes $N = \{1, ..., n\}$. In our context, each individual agent (user) is represented by one node. Further, for each agent i, there is a vector $g_i = \{g_{i,1}, ..., g_{i,i-1}, g_{i,i+1}, ..., g_{i,n}\}$, such that each $g_{i,j} \in \{0, 1\}$ indicates whether or not a direct link exists between nodes i and j. These links can be interpreted as friendships between users in our context. A network g therefore is defined by the set of vectors g_i, $i \in N$. All friendships have the same "strength" in this framework, as they are depicted as purely binary relationships.

3.3.1. Creation and Severance of Individual Links

Assume that each link g_i, j between two users or players i and j is undirected, i.e. both can benefit. To form a new link, both players have to consent to the formation. One player (the "sender") initiates the link formation process through an invitation while the other (the "receiver") decides whether to accept, ignore or decline the invitation. If the receiver accepts the invitation, a new link is formed and $g_{i,j}$ becomes 1. We denote this change in the network structure by $g' = g + g_{i,j}$. If the receiver declines or ignores the invitation, the structure of the network does not change.

Conversely, each user inside a friendship can sever an existing link unilaterally. If user i or user j cancels their friendship, then $g_{i,j}$ becomes zero. Analogously to above, we denote this change in the network structure as $g' = g - g_{i,j}$.

The utility that a user obtains from the network depends on its structure, therefore $u_i = u_i(g)$. It appears sensible to assume that the formation of a new friendship is associated with costs, which may vary with the characteristics of the individuals involved: The sender must first locate the receiver, then click on a button and compose a message to send out a friendship invitation, and finally he loses one free slot among his friendships if the receiver accepts the invitation. Once a friendship offer has been extended, the search costs are sunk and only the costs of the free slot remain. In our context, all these costs are most likely negligible; ignoring them would not change the following analysis. There is a certain asymmetry here: The receiver only has to click on one button, whether he accepts or rejects a friendship (she has to take no action to ignore the invitation) and she looses one free friendship slot if she accepts.

As the formation of new friendships requires mutual consent, the following two conditions have to hold at the same time for a new friendship to be created, assuming that i is the sender and j is the receiver, and denoting their respective costs of inviting/accepting as $\kappa_{i/j}$:

$$u_i(g + g_{i,j}) - \kappa_i > u_i(g) \Leftrightarrow \Delta u_i > \kappa_i \tag{3.1}$$

$$u_j(g + g_{i,j}) - \kappa_j > u_j(g) \Leftrightarrow \Delta u_j > \kappa_j \tag{3.2}$$

The way the Network is set up, additional links are close to purely beneficial. They allow users to access more music and to make better use of additional play-lists. As indicated above, users in the Network are mainly interested in music consumption, therefore these considerations should be their main priority. On the other hand, each additional friendship costs next to nothing as long as users are not close to the upper limit of friends, which is given for the period of time that we study.

Analogously, for user i to be willing to sever an existing link, denoting the costs of severance as λ_i the following condition must hold:

$$u_i(g - g_{i,j}) - \lambda_i > u_i(g) \Leftrightarrow \Delta u_i > \lambda_i \tag{3.3}$$

For the same reasons as stated above, we would expect Δu_i to be negative in the case of link separation, unless the number of friends is 150 (or close to the upper limit and users anticipate reaching it with certainty). As the costs of link separation λ_i are non-negative, in this simple setting users are expected not to separate links.

3.3.2. Provision of Club Goods in Network

One of the main features of the Network is that users provide different kinds of effort, uploading music or compiling play-lists, that are potentially useful both to themselves and to their friends. We treat this in the simplest imaginable form: Consider that user i can exert effort e_i at cost $c(e_i)$. We denote the set of user i's first degree friends as J_i, each of whom also exerts effort. Expanding the utility function of user i above to incorporate these considerations, it now becomes $u_i(g, e_i) = u_i(e_i, \{e_j\}_{j \in J_i}) - c(e_i)$. Clearly, the optimal level of effort of i is determined by

$$\frac{\partial u_i(e_i, \{e_j\}_{j \in J_i})}{\partial e_i} = c'(e_i) \tag{3.4}$$

Therefore, if the effort of user i and his friend user j are substitutes (i.e. the cross-partials are negative) and the costs of effort provision are non-negative, the optimal effort of i should decrease with the effort that his first degree friends provide. If in addition the utility of user i increases in the effort of his friend user j, i.e. $\frac{\partial u_i(e_i, \{e_j\}_{j \in J_i})}{\partial e_j} > 0$, we can talk of a local public good.[9] But how likely is it that these assumptions are satisfied in the setting we encounter?

1) Utility increase in effort provision of friends: More songs provided by his friends gives a user more songs to choose from and likewise more playlists give him a combination of more choice and better choices. Standard

[9] Note that in the Network, we can distinguish two dimensions of effort, provision and organization of information, which are complements.

theory usually states that more choice is a good thing, recent advances in behavioral economics non-withstanding. As the users of the Network display many characteristics of collectors and music connoisseurs we find arguments to the opposite unconvincing.

2) Negative cross partials: This point is less clear-cut. If someone uses the service to listen to a given set of songs that he has decided upon ex ante, then there is no substitute for these songs. On the other hand, if users want to create a mix of music to entertain them, then they will value variation above all else and a song uploaded by a friend is an almost perfect substitute to a song uploaded myself. We do observe users listening to their friends' music frequently, as one can glean from the analysis below, therefore we find this assumption to be equally justified.

It is straight-forward to derive the expectation of free-riding by users from these considerations. In fact, in a similar setup, Bramoullé and Kranton (2007) find a unique stable equilibrium with complete specialization, i.e. the agents can be cleanly separated into two groups, one of which provides effort, the other of which free-rides entirely.

There are two straightforward reasons why this is not necessarily so clear cut in the Network we are considering. For one, users may not take the link-structure as given, but can instead strategically form friendships and sever existing links. Therefore social norms may be formed and upheld that free-riding can and is to be punished by severing links with free riders, which we refer to as negative reciprocity. In addition, users may only enter into new friendships with others who have already demonstrated that they are good citizens (strategic link formation), i.e. who have provided measurable effort up front.

The second reason why free-riding does not necessarily have to occur is the fact that many of the peer-groups within the network are relatively small. Even the most active group of users distinguished above, those who have logged into the network in the course of the past 7 days, only has around 4 friends on average. In such a close (if still potentially anonymous) setting, one may expect that social forces such as shame or guilt, combined with the relative ease of monitoring a small number of connections, may prevent users from engaging in free-riding behavior. Experimental evidence shows, though, that even in small groups the play in repeated public-good games converges neither to a pure free-riding equilibrium (the rational equilibrium), nor to purely cooperative behavior (the efficient outcome), but instead behavior is mixed, with play converging towards a situation in between the two extremes. See for example the seminal contribution of Andreoni (1988). We will simply take these competing hypotheses to the data.

From all of the considerations above, we derive the following hypotheses:

Hypothesis 1a - Strategic Linking Behavior regarding Invitations: Users that are invited to join a friendship have provided higher efforts on average than the peer-group of active users.

Hypothesis 1b - Strategic Linking Behavior regarding Acceptance: For friendship invitations that are accepted, the inviter has provided higher efforts on average than her peer-group. Especially, users whose friendship invitations have been accepted have provided significantly more effort than users whose friendship invitations have been denied or ignored.

Note that for the formation of new friendships, both users have to agree, therefore the behavior and characteristics of both parties are relevant for the formation. As opposed to this, for the severance of friendships it is sufficient for one of the users to make the decision. This is reflected in the following hypothesis, which will be more closely specified further below:

Hypothesis 2 - Negative Reciprocity: Friendships from which one of the users derives little utility are more likely to be severed. Especially users who display behavior that resembles free riding are more likely to be excluded from friendships.

The hypotheses concerning free riding behavior are very straightforward:

Hypothesis 3a - Free Riding for Information Provision: Users upload fewer songs themselves, the more songs are available to them through their friends.

Hypothesis 3b - Free Riding for Information Organization: Users generate fewer play-lists themselves, the more play-lists are available to them through their friends.

While our first three hypotheses focus on free-riding and negatively reciprocal punishment behavior, conversely one could also expect a social norm to exist that rewards desirable behavior or effort provision in some form. The design of many websites takes exactly this form of reward into account: For example, in many popular specialized internet-forums, users are awarded ascending titles based on the number of contributions they have written. Similarly, Wikipedia-contributors can access statistics regarding the number and extent of their contributions.

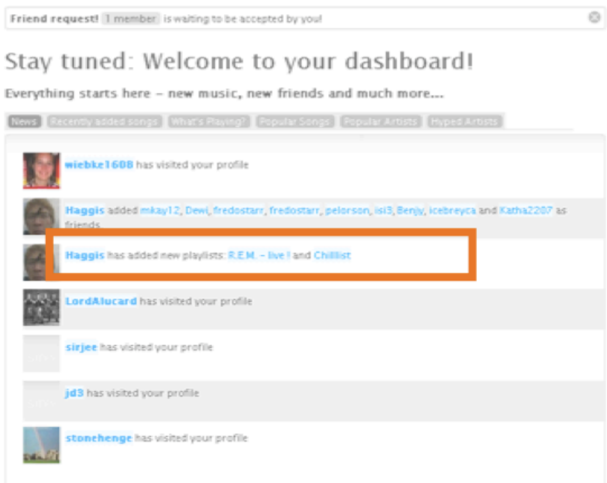

Figure 3.2.: Greeting screen with information of friends' activity.

Also the design of the Network we study has a built in perquisite of this form. Whenever a user logs onto the service, he is greeted by a welcoming-screen. Placed centrally on this screen is a message box that informs the user whenever one of their friends has supplied additional play-lists or music, as depicted in **Figure 3.2**. Confronted with this information regarding the beneficial activities, users might want to positively reciprocate by providing additional effort of their own, as for example suggested by Fehr and Gächter (2000). The potential reasons for this kind of behavior are manifold: They range from competitive instincts (who has the most songs in his library), over imitation to gratitude and reciprocation. Irrespective of the true motivation, we can subsume each of these under the following two hypotheses:

Hypothesis 4a - Positive Reciprocity for Information Provision: Users respond in kind whenever their friends provide additional information/upload new songs, i.e. users are more likely to upload new songs if their friends have recently uploaded new songs.

Hypothesis 4b - Positive Reciprocity for Information Organization: Users respond in kind whenever their friends provide additional organization of information information/create new play-lists, i.e. users are more likely to compile new play-lists if their friends have recently generated new play-lists.

3.4. Empirical Analysis

3.4.1. Data Description

We use two different data sources to approach the questions that we wish to study. For the questions related to link formation and severance, we mainly use data covering every change of user friendship status between July 10th and November 11th 2008. From this we create two different data-sets. To test Hypotheses 1a and 1b, we look at all friendship invitations that were issued during this time, as well as what happened to them (acceptance, declination, revocation, nothing). We observe 3,657 friendship invitations in this time-span. Importantly for our topic of interest, we also observe certain user characteristics for those involved in these exchanges. In order to estimate the probability for a friendship being severed for Hypothesis 2, we limit our analysis to those invitations that were accepted in this time-span and which we observe for at least 50 days. This limits our number of observations to 1,52. We call this data set **A**.

For the remaining hypotheses on free riding and positive reciprocity, we make use of a panel data set. We have 10 weekly cross-sections of user data in which we observe various characteristics and behavior variables as described below. The data cover the period from September 14th until November 18th 2008. Due to attrition and new arrivals, this results in an unbalanced panel with ca. 25,000 users and an average of 8.1 observations per user. We will refer to these as data set **B**.

In the following, we briefly describe the variables used in our study and indicate in which of the data sets they are available to us:

> Days Online (A,B) - This count variable logs each day that a user logged into her account or used one of the functions on the Network. The value of this variable is the aggregated number of days that a user has visited the site. Note that this variable does not measure how many actions a user has performed on a given day or how much time he has spent using the service.
>
> Days since Registration (A,B) - This variable describes how many days ago a user registered with the service.
>
> Last Login (A,B) - Measured in days, this variable denotes the time since the last activity of the user. An activity is a login or any other action within the user interface.
>
> Friends (A,B) - The number of currently active friendships a user is involved in. There is a limit of 150 friendships that a user can enter into, but this limit was binding for only one user at the time of our analysis.[10]
>
> Music Uploaded (A,B) - This variable states the number of uploaded

[10]Due to a technical glitch in the system, this user was actually able to enter into more than 150 friendships.

songs by a specific user. These songs are collected in the users own music library, which is accessible for the user's friends.

Songs Listened (A,B) - This is the total number of songs the user has listened to on the Network. A limitation to this variable is that it includes both songs that were skipped as well as songs that were listened to in full length. Nevertheless, assuming that users derive utility from listening to music through the server, this is one of the better proxies for user utility available to us.

Own Play-lists (A,B) - The number of play-lists that a user has created and that exist at the current point of time, i.e. if a user deletes a play-list, this variable declines by one.

Play-lists Friends (B) - The total number of current play-lists of all of a user's friends.

Common Friends (A) - All users that are friends of both the sender and the receiver of a friendship invitation at the given point of time are counted in this variable, illustrating the overlap of senders' and receivers' friends.

Common Music (A) - This variable counts songs that are owned by both sender and receiver. Important to note is that this variable is not able to completely capture the overlap of the both music libraries, since some songs may vary in quality or length and are thus tagged with different internal IDs by the system.

Songs from Friend's library (A) - This counts the number of songs from the other users library that a user has actually listened to.

3.4.2. Strategic Linking Behavior and Negative Reciprocity

One criticism that one could easily point at our project is that users do not act out of rational objectives, i.e. that they interact with people they consider friends totally irrespective of their effort provision. The aim of the first hypothesis is to if not repudiate this argument then at least to put it into perspective. **Table 3.2** below shows the average characteristics of four different groups: First, those who received friendship invitations throughout the period we observe. We partition the senders of these invitations into two subgroups, those whose friendship was accepted and those whose friendship was declined or left pending for more than 30 days.[11] We compare the characteristics of these groups to the group of those that have logged in within the past 10 days, the most restrictive "conventional" activity measure.

The results are quite striking. People who have received an invitation are on average by far the most active group. They have uploaded far more music (3 times the average of the second most active group), provided more play-lists and are also more active concerning the social network functions. Both mean- and median-equality tests are rejected at the 1% significance level for each of these variables. Therefore we cannot reject **Hypothesis 1a** and consider the evidence in support of the hypothesis to be strong. The users in the network appear to be guided by rational considerations in deciding whom to invite to form a friendship.

Similar arguments hold with respect to **Hypothesis 1b**. Those users whose invitation was accepted have provided twice as much music as their rejected (or ignored) counterparts, almost three times as many play-lists and have sent twice as many messages or comments on average. Again, both mean- and median-tests show that the differences are significant at the 1% level. So the evidence is strongly in favor of the hypothesis. For the more general point, we also conclude that users do behave in a way that appears to be driven by

[11] The date was chosen as more than 95% of acceptances were issued within 28 days, but the results are highly robust to moving the cutoff date.

	Invitees	Inviters Accepted	Inviters Not Accepted[a]	Act10 users[b]
Own music	1936.02	667.30	316.28	368.71
	(5416.59)	(1879.08)	(1326.79)	(1299.06)
Own play-lists	40.65	20.95	7.79	11.12
	(126.97)	(90.41)	(37.37)	(47.00)
Friends	38.07	21.5	20.61	14.48
	(39.76)	(27.50)	(26.69)	(27.93)
Messages sent	36.41	10.40	5.61	11.47
	(150.83)	(62.61)	(33.44)	(88.03)
Comments left	2.82	1.17	.59	1.06
	(7.82)	(4.16)	(2.88)	(4.78)
obs.	3657	2,076	930	602

[a] An invitation was considered non-accepted if it was outright declined or if it was left pending for more than 30 days.

[b] Here, we exclude those users that have registered less than 10 days ago, which explains the differences to table 1.

Table 3.2.: The average characteristics of various groups of users: invitees, accepted and non-accepted inviters and the most active group of the Network's users according to the most widely used definition, (means reported, std. deviations in parentheses)

rational (utility maximizing) motives.

Next, we focus on the unilateral decision to sever existing friendships. We distinguish between two parties, the person who severed a friendship ("Severer"), and the person who passively had to endure the severance ("Severed"). Even a cursory glance at the descriptives displays a very clear pattern. In **Table 3.3**, we present the average characteristics of a Severer, a Severed and an "average" user in the sub-sample. Comparing the latter to the other measures of user characteristics presented above shows us that the new friendships we observe were formed between relatively active and involved users.

The comparison between Severer and Severed is very stark: From a static viewpoint, the Severer had provided 22-times the music (4472.15 vs. 206.14) and 48-times the number of play-lists (358.62 vs. 8.47) at the time that the friendship was formed. Looking at their individual behavior during the time that the friendship lasted, the Severer uploaded 110-times as many songs per day (33.71 vs. 0.28) and 12-times as many play-lists (0.65 vs. 0.051) as the Severed. Finally, the Severer was far less likely to listen to a song of the Severed than vice-versa (0.02 vs. 1.29 songs of the other listened to per day). The small, yet positive, number of songs that the Severer listens to from the Severed's music library is a reminder that the decision to end the friendship is costly: You lose access to those songs that the other person has provided. In a static world, which takes the utility derived from a friendship as given, this appears not to be rational, unless a user derives an immediate satisfaction from link severance ("revenge" for example). On the other hand, it is relatively easy to imagine a dynamic game in which an equilibrium exists, in which punishment counteracts free riding. Both concepts have in common that free riding is the cause of dis-utility to the other party. To restate our second hypothesis in accordance with this argument:

Hypothesis 2 - Negative Reciprocity: Friendships in which one of the users displays behavior resembling free-riding are more likely to be severed.

Importantly for our analysis, no user reached the maximum number of friends (150) during the time that we observe, the maximum number of friends that was reached in the severed friendships was 125.

From the descriptive statistics alone, it is hard to disentangle whether it is the sheer amount of effort provided by a user or the additional provision of effort over time that the other party values. In order to shed more light on this question, we consider the following discrete decision model. Let us call the probability that a given friendship is severed $p(x) \equiv P(y = 1|x)$, where x is a vector that captures both parties' characteristics. Using the latent variable approach, we estimate logit-models, see **Tables A.1 and A.2** in the Appendix for the exact results.

In model 1a, we only include the "static" characteristics, i.e. the effort provided by both users at the beginning of their friendship (here again, we talk of a severer and a severed user - if the friendship was not severed, then these two roles were assigned to the users in a friendship randomly). We would expect

	Severer	Severed	Average User [a]
Own music beginning	4472.15 (3111.36)	206.14 (987.39)	885.84 (1335.54)
Average songs added per day	33.71 (34.44)	0.28 (1.03)	4.08 (11.05)
Own play-lists beginning	358.62 (317.54)	8.47 (36.25)	31.52 (71.19)
Average play-lists added per day	0.65 (1.80)	0.051 (0.11)	0.29 (0.84)
Friends beginning	12.82 (29.15)	15.61 (22.90)	28.96 (22.68)
Friends end[b]	22.27 (35.36)	17.63 (23.24)	34.48 (24.34)
Number of friend's songs listened per day	.020536 (0.13)	1.29 (3.65)	0.15 (0.74)
obs.	114	114	3,042

[a] Average of users in the entire sample of new friendships being formed.
[b] The maximum number of friends observed among severed friendships was 125, well below the limit of 150. For non-severed friendships, we report the number after 50 days.

Table 3.3.: Mean characteristics of users involved in severed friendships: The severers of friendships compared to those who were severed and the average of all users involved in "new" friendships.)

that the effort provided by the Severed should lower the probability that the friendship will be ended, as it raises the costs of the Severer (more information that he can no longer access), and the effect does have the expected sign and is (though weakly) significant for the provision of music, while the provision of play-lists is not significant. Interestingly, the effort provision of the Severer is significant and increases the probability of severance along both effort dimensions - this means that users that have contributed a lot of effort themselves are more likely to end a friendship (while they are also more likely to encounter free riders, according to the analytical framework presented above).

In model 2a, we only include the "dynamic" characteristics, i.e. the number of songs and play-lists that the individual users have added over time on average. We encounter the same pattern: The more music the Severer adds, the more likely he is to end a friendship, while adding more music makes it less likely that a user will have his friendship ended by the other party. The effect of play-lists added by the Severer is more interesting. Adding more play-lists makes it much *less* likely that a user will end a friendship. On second glance, this is rather intuitive: When adding a play-list a user organizes the information available to herself for easier access. This may (or may not) include information provided by the other user. In any case, this makes the information more valuable (as it is more accessible) and therefore raises the costs of ending a friendship.

In model 3a, we include both the static and dynamic characteristics. Encouragingly, the effects retain their signs and significance levels.

In models 1b and 2b, we add those variables that capture free riding in the potentially most intuitively appealing way, by including the number of songs of the other person that the respective users have listened to per day. Model 1b only includes these two variables and, as one would expect, songs listened by the severer decrease the probability that a friendship is severed. Further, the number of songs listened per day by the severed highly significantly increase the probability that a friendship will be severed. This resembles the most intuitive description of free-riding, but it is noteworthy that the severer does not know that his music library was accessed by the other party. In model 2b we add all explanatory variables used in the previous models and the effect that this intuitive form of free-riding strongly increases the probability of a friendship being severed persists.

To conclude the section, let us briefly summarize our findings. There are strong indications that users make their link-formation decision contingent on the utility they expect to derive from a given friendship. Invitations are sent out to very active individuals, and invitations by users who have provided little effort themselves are far more likely to be rejected. Regarding the severance of links, we find that users are willing to incur severance costs (not being able to access the other's information) in ca. 7.5% of friendships in the interval of time that we observe. Patterns and regularities, as well as our logit regression results, suggest that this behavior is caused by the desire to punish the other users' free riding. This corresponds with experimental findings especially in the literature on the ultimatum game. Interestingly, this punishment behavior can itself be interpreted as a form of public-good provision, e.g., enforcement of social norms.

What remains now is to show how prevalent free-riding is in this community, given the existence of the described control- and punishment mechanisms on an individual friendship level.

3.4.3. Free Riding and Positive Reciprocity

As discussed above, we consider a user to be free-riding if she provides less effort ceteris paribus given that her friends have exerted themselves more. A naive approach to this problem therefore would be to simply regress all users' levels of uploaded songs and generated play-lists onto the levels of their friends in a cross-sectional approach. At second glance, though, this approach is clearly not conductive to answering the question due to a reverse causality issue. Assume that one finds that people whose friends provide a lot of effort provide a lot of effort themselves. One can clearly imagine a story in which people who have provided a lot of effort in the past attract more friends, especially considering our findings on strategic linking behavior.

To circumvent this issue, we therefore don't consider the total levels of effort that users provide as a dependent variable. Instead, we look at their weekly efforts, i.e. how many songs/play-lists they add in a given week, given the explanatory variables and different sets of controls from which we obtain the following regression equations:

$$e_{it} = \beta_1 X_{it} + \beta_2 Z_{it} + u_{it} \tag{3.5}$$

where e_{it} is the effort an individual exerts, X_{it} describes the explanatory variables and Z_{it} designates the controls. Clearly, there may be unobserved individual effects that are correlated with the individual error terms, e.g. some individuals may be compulsive collectors of music, while others are casual listeners only. In order to capture these differences, we estimate a user fixed-effects model, and adjust the standard errors for potential user-cluster effects.

We estimate three different models (with two different samples) for information provision and organization respectively, in order to shed some light on hypotheses 3 and 4. In each specification, X is composed of three variables:

The determinants for the effort that a user provides for **information provision** in a given week should be determined by a) the total amount of music that his friends have made available to him, b) the total amount of music that he himself has uploaded in the past and c) the amount of music that his friends have added in the recent past. For c), we use the additional music friends have uploaded in the course of the past week - if anything, this should under-estimate the effect we are trying to find. Relating these variables to our hypotheses we would expect the following: a) should have a negative influence on effort provision (free riding, hypothesis 3a), b) should have a negative effect on effort provision due to decreasing marginal utility from effort and c) should have a positive influence on effort provision (positive reciprocity, hypothesis 4a).

The three models differ by the controls that we add. Model 1 is bare, i.e. without additional controls. In model 2 we add a control for the user's time since registration by adding a dummy that takes the value 1 if the user has been registered for longer than 60 days (experienced dummy). Further we add a control for time. In order to control for users whose main motivation to join the network is mainly social, we add controls for the amount of public comments that users have left and private messages they have written. In model 3 we give credit to the consideration that the two kinds of efforts are complementary. One would expect users to add more music if their existing music is well organized and vice versa. Therefore in model 3 we add controls for the number of own play-lists and number of friends' play-lists available to a user. We estimate each of the models once for the entire population of registered users and once for the sub-sample of active users only.

Analogously, for **information organization** efforts in a given week, the explanatory variables in X are a) the total number of play-lists that a user's friends have provided, b) the total number of play-lists that he himself has compiled and c) the number of play-lists that his friends have added in the past week. From our hypotheses, we would expect the same signs as above. Again, model 1 is bare, model 2 adds controls for time and social interactions and model 3 takes into account the amount of music that the user and her friends have respectively uploaded.

For the detailed regression results, please refer to **Tables A.3 through A.6** in the Appendix. We find very weak evidence for hypothesis 3a, free riding with regard to information provision. Users upload weakly significantly less music only for the sub-sample of active users in model 3, i.e. taking the effects of the complementary effort provision into account. On the other hand, the evidence for positive reciprocity with regard to information provision (Hypothesis 4a) is comparatively robust. In each specification of the model there is a small but significant and positive reaction to other users' providing more music. Some other observations for the regressions regarding music files added are highly interesting. For one, the level effect of the own music uploaded is never significant - this could either be due to the fact that the marginal utility from additional music is non-decreasing, i.e. now matter how much music is already available, more music is yet better. Or there is so much music available freely in any case, that additional own music files from the beginning on have little value to users.

The regressions for the additional provision of play-lists **(Tables A.5 and A.6)** resemble our assumptions rather closely. There is a significant negative level effect for the amount of own play-lists already compiled, which resembles initially positive and decreasing marginal utility of organization of information. Model 3 shows that the more music a user himself provides, the more likely he or she is to exert more effort in organization. The effect of friends' music on the other hand is surprisingly significant and negative (if economically not very meaningful). Again, this might resemble the fact that more music in questionably accessible libraries may actually reduce the utility a user derives from the service because it makes things harder to find. Regarding our central hypotheses, we find robust evidence for hypothesis 4b, users react to their friends' organization of information in a positively reciprocal manner. Again, the evidence for free riding is limited to one of the sub-samples and only appears when controlling for the complementary level of information provision in model 3. Finally note that the absolute value of the coefficient of the positively reciprocal effect is twice the size of the free-riding effect - therefore the net effect of a friend providing more music is *positive* in the short run.

3.5. Conclusion and Outlook

In this article, we attempted to empirically assess the rationality of user behavior, analyzing unique data from a real social online-network. The major advantage of this data is that we observe both individual user characteristics as well as the structure of the entire network which allows us to study issues that were not easily approachable up until now empirically.

We find that despite mixed experimental evidence, the existing theory on link-formation allows straightforward and useful predictions: Users systematically pursue more valuable "friends" in the process of link formation. On the other hand, individuals are willing to sever at first glance purely beneficial links in 7.5% of the friendships that we observe. There is substantial evidence that this is linked to punishment-behavior when users encounter free riding, i.e. a direct form of negative reciprocity that in the past has been most prominently encountered in ultimatum games. Viewed statically, this kind of behavior would not be predicted by network theory (or economic theory in general). Interestingly, this can in itself be construed as a form of public good provision, as it may help to enforce social norms, from which users in general benefit. Finally, we focused on whether free-riding is detectable, or whether on the contrary we observe a form of positively reciprocal behavior in which individual users respond to their friend's additional effort provision in kind. There is clear evidence for the latter, with the positive effects of reciprocal behavior overcompensating the free-riding effects in the short run.

We believe that our research raises some interesting theoretical questions.

Evidently, there are interactive effects between the formation of links, the provision of club goods within a network and unilateral link-severance, which may influence the prevalence of free-riding within a social network. These effects have, to our knowledge, not been captured in theoretical network models up until now, even though they may yield predictions that are far closer to the empirically observed facts. In addition to the individual behavior, it would be interesting to generate predictions concerning network formation and equilibrium network structure in an in this sense richer context.

For the specific network, there are a number of issues that we want to address in future work with additional data. The questions include, but are not limited to: Is the punishment mechanism that we observe effective, i.e. do individual users, who have been "kicked out" of friendships change their behavior? What are the effects of the size of neighborhoods (or the number of friends of an individual user) on his free-riding behavior - e.g. are members of larger neighborhoods more likely to exhibit this kind of behavior, as potentially monitoring costs increase and punishment is less likely? Finally, it would be wonderful to see whether our findings hold true for comparable social networks, if similar data becomes available in the future.

In closing, we would like to restate that we are well aware of the drawbacks that empirical work in this field suffers from in general. We do not observe and cannot perfectly control for individual users' actual motivation in their actions. There may be substantial and unpredictable differences in their utility functions, "network savvy" and computer literacy, which may systematically bias our results. While experimental studies are not nearly as prone to these problems, they are subject to certain shortcomings of their own. We hope that our study is viewed as a strongly complementary building block that enriches our understanding of the behavior of individuals in complex social network settings.

4. Trust and Investment - A Theoretical and Empirical Assessment

4.1. Introduction

4.1.1. Motivation

In the last couple of years, the automotive sector, one of the most innovative and important industries in developed countries, has witnessed unprecedented turbulence. In June 2009 General Motors, the second biggest carmaker in the world, filed for chapter 11 bankruptcy proceedings. It appears that more than just being strangled by mounting pension obligations and being undermined by rising fuel-costs, the former industry juggernaut has been out-manouvered by more innovative and design-savvy competitors. Since then, the financial crisis has made consumers even more loth to make extensive investments into their mobility.[1] Increasingly harsh competition for a slowing demand has eroded industry margins, placing an ever greater premium on companies that produce efficiently and are able to differentiate themselves from others through their high level of quality and innovative products; but to an ever greater degree, it is upstream suppliers' know-how and efforts that are responsible for successful differentiation.

As a result of increased outsourcing efforts, carmakers contribute only a minor share of innovation efforts and the total value added of their product -

[1]Despite the fact that governments throughout the world have responded decisively and expensively, propping up ailing carmakers by guaranteeing loans and providing liquidity on the one hand, and giving consumers additional incentives to replace their old cars through variations of the "cash for clunkers" program.

in some cases they are almost reduced to pure assemblers. Upstream suppliers are responsible for much of the ground-breaking basic research, which is then adapted to the specific needs of individual car models. Nevertheless, the structure of the market places the car manufacturers in the superior bargaining position vis-à-vis their suppliers. This leads to an inherent dilemma between attaining the desired quality levels and short-term rent-extraction/profit generation. We demonstrate that *trust* can play a central role in such a hold-up situation in a theoretical model, whose predictions we then test using a unique data-set.

To be able to do this, we contribute to the understanding of trust: While there have been significant advances in economic research on this topic in the past decade, many of them share what can be considered a drawback: Both experiments as well as empirical approaches focus on trust as a characteristic of subjects, either investigating their general attitudes towards others ("Do you think in general others can be trusted?") or analyzing their willingness to contribute funds to more or less anonymous players in lab settings. In contrast to this, we define an upstream supplier's trust in the downstream firm as his expectation (the subjective probability) of whether the latter is going to exploit a holdup scenario to extract rents. This is closer to the sociological and colloquial reading of the term, in which in a specific setting A trusts B to do something (or refrain from doing something), which is in B's realm of influence. With this understanding, trust can alleviate the under-investment problem. We further show that, as opposed to many results in the literature on relational contracting, trust and competition between upstream suppliers induced by the downstream firm are not mutually exclusive. In fact, competition can be a substitute for other forms of rent extraction and in this case stronger competition induced by the manufacturer would be associated with higher levels of suppliers' trust.

We are able to approach the subject of trust empirically using a unique data-set, collected from an online-survey of suppliers and manufacturers in the German automotive industry in 2007-2008. The measures of trust that we generate are relationship-specific – instead of subject specific – and we analyze the determinants of suppliers' trust in the manufacturer by linking them to reported past behavior. We show that higher levels of suppliers' trust lead to higher relationship-specific investment proxied by failure-rates of supplied parts. Finally we show that more intense supplier competition is not mutually exclusive with trust in the downstream firm.

The remainder of the article is organized as follows. After briefly outlining the related literature, we develop a very simple model in Section 2, from which we derive hypotheses on the effects of trust on vertical relationships. In Section 3, we first introduce the study that provided the data on which we base our empirical analysis, an in-depth survey investigation into the structure of the German automotive industry. We present potential measures of trust and try to carefully evaluate what they capture. At the center of our empirical analysis, we analyze how trust between manufacturers and suppliers is related to two important questions: Sourcing decisions and supplier (under-)investment. Finally, Section 4 concludes and raises some new research questions resulting from our findings.

4.1.2. Related Literature and Contribution

Interactions between suppliers and OEMs[2] in the automobile industry are notoriously complex, fraught with moral hazard and hold-up problems. Nevertheless, both parties regularly invest substantial amounts of time, know-how and money into specific relationships. Apparent puzzles like this have piqued the interest of economists for quite some time—perhaps the most prominent explanation approaches can be subsumed under the headings of property rights theory (applied mainly to hold-up problems) and contract theory (applied mainly to asymmetric information and moral hazard).[3] Beyond that, it has been well

[2]We use the term *OEM*, i.e. original equipment manufacturer, for the downstream automobile producer. We will refer to the upstream firms simply as *suppliers*.

[3]Due to the more applied focus of our study, we refrain from delving deeply into the intricacies of the literature on the holdup-problem. Instead, we refer to the seminal Hart and Moore (1988), Grossman and Hart (1986) and Hart and Moore (1990), as well as the more recent Hart and Moore (2007) and

established that in settings like these, relational (or informal) contracts can play an important role in governing relationships. As opposed to formal contracts, which are linked to outcomes verifiable by third parties and courts, the term relational contract refers to *self-enforcing*, often implicit agreements "sealed with a handshake".

Theoretical Approaches

There is a rich theory on relational contracting in different contexts,[4] beginning with Bull (1987) who provides the original repeated games-framework for employment relationships. Also for employment, Baker et al. (1994) demonstrate how the combination of formal and relational contracts can lead to better results than either instrument could achieve alone. Interestingly, there can be a substitutive (if either works almost perfectly) or complementary relationship between the two. Baker et al. (1999) analyze the informal delegation of decision rights within hierarchies under different informational settings.

More recently, research has focused on more generic settings, searching for optimal contract design. Levin (2003) finds that while under moral hazard optimal relational contracts exist and are relatively simple, under hidden information cases arise in which agents do not respond to the incentives provided therein at all. Calzolari and Spagnolo (2009) further extend this scenario: The relationship between a principal and an agent interacting repeatedly can suffer from both moral hazard and hidden information - but the principal has a further tool available to him by being able to select from various competing agents (screening) who are able to collude. They find the intuitively appealing result that in cases in which non-contractible factors contribute more to the principal's payoff, the best package of instruments will rely more heavily on relational contracts with a smaller set of agents. Vice versa, in "simpler" settings in which the most important issues are contractible, the principal will rely on

Hart (2008), the latter two with many further references.

[4] We refer to MacLeod (2007) for a careful survey of the literature on relational contracting, with a special focus on the effects of the quality of legals systems.

a more competitive setting amongst agents. Brown et al. (2004) carry out an experimental study in order to be able to control for the level of enforceability of contracts. They find that as enforcement becomes more effective, the original long-term rent-sharing relationships are replaced with short-term arms-length agreements, very much in line with the theoretical results described above. This can be taken as evidence of a substitutive relationship between formal and relational contracts in their specific setting.

Entering into a business relationship without being able to resort to legal means of enforcement would probably be called a form of **"trust"** colloquially. The models described above have in common that relational contracts and competition are in a certain sense polar concepts—you either rely on handshakes or you prefer arms-length market interactions. We will demonstrate in our analytical model (as well as the empirical part of our study) that trust and competition can actually be *associated* with each other in a hold-up setting, if competition is used as a "legitimate" mechanism of extracting rent from the individual, instead of "illegitimately" extracting rents through superior ex-post bargaining power. As the concept of trust that we propose is not congruent with relational contracts, this result can be seen as complementary to the existing literature.

Empirical Studies

Trust has for some time drawn considerable attention and scrutiny from experimental economists.[5] Yet researchers in applied micro-economics and industrial organization have been rather cautious about using this term, even actively trying to avoid it (see, e.g. the discussion in MacLeod (2007)).[6] On the other hand, empirical researchers in the areas of macroeconomics and growth have been less reticent in this regard, so that some empirical strategies do already exist.

[5] See Fehr (2009) for a sweeping overview of the experimental and neuro-economic literature.

[6] For a careful survey of the development of the term "culture" in economics and the effects of culture on economic outcomes, see Guiso et al. (2006).

As a basis for many studies, the answers to the following question from the World Values Survey has been used: "Generally speaking, would you say that most people can be trusted or that you have to be very careful in dealing with people?" While one may doubt the power of this construct at first glance,[7] it has been used frequently to obtain results. The basic hypothesis of La Porta et al. (1997) is that trust is an integral requirement for the functioning of larger organizations in which the likelihood of repeated interactions is relatively small and thereby the established mechanisms for ensuring cooperative behavior are less effective. In a cross-country study they try to establish that populations in which higher levels of trust are prevalent should foster more effective governance as well as relatively larger firms. Aghion et al. (2008) perform an international comparative study in which they scrutinize the connection between levels of social capital (or trust/distrust) in populations with the amount of existing state regulation as well as the demand for it. The basic intuition is that a lack of civic mindedness in one's fellow citizens may lead to a stronger desire for the state to regulate interactions. They find very strong evidence for this, even for societies in which the government itself is plagued by corruption. Therefore it appears that trust and regulations are to some extent substitutes. In contrast to most other articles cited so far, Butler et al. (2009) study the effects of trust on individual's economic outcomes instead of aggregate economic performance. They use the European Social Survey as well as experimental evidence to argue that a medium amount of trust may be optimal for individuals: With too little trust, too many opportunities for beneficial interactions are missed, with too much trust the danger of being taken advantage of becomes too great.

Guiso et al. (2009) use a slightly more concrete measure, the trust that citizens of a given country in Europe have for citizens of another. They find that the levels of trust are explained in part by characteristics such as the distance between countries, but also by factors such as sociological and genetic closeness and common history. They find that less trust in the citizens of a

[7]See Sapienza et al. (2007) for an experimental study on the merits of this measure and a discussion of the previous literature.

country is associated to significantly lower aggregate trade and investment. In a second study, Guiso et al. (2004) suggest that the different characteristics of Italian regions lead their citizens to develop different levels of social capital. They then show that in high social capital/trust areas people are more prone to invest in stocks instead of holding cash reserves and have easier access to bank credit. The effect is mitigated by levels of education. Along similar lines, Guiso et al. (2005) find that individuals who display higher levels of trust buy more risky assets relative to their wealth. They counter a common criticism by controlling for risk- and ambiguity-aversion, which does not make their original result disappear.

Finally, Bottazzi et al. (2009) study the willingness of venture capitalists to perform non-contractible services in a micro-economic environment. In particular, they analyze the influence of more effective legal systems in this context. Both in their theoretical model and their empirical analysis of a data-set with European venture-capital deals they find that a more efficient legal system has two effects. On the one hand it is complementary to trust, in the sense that it makes venture capitalists more willing to grant non-contractible support; but on the other hand, they require more protection for the case of failure of the venture.[8]

We believe that the empirical part of our study adds a new angle to the way economists think about trust. We define trust not as a characteristic specific to individuals. Instead, we interpret trust as the expectations of one party towards future behavior of the other, based on relationship-specific observable and unobservable characteristics, including the joint history. Based on this approach, we attempt to analyze how higher or lower levels of trust in a downstream firm by an upstream firm affect upstream investment levels as well as the choice of contractual setting in the German automobile industry.

[8]In the working paper on the same data, Bottazzi et al. (2007) also show that higher scores on the Eurobarometer measure of trust between nations are associated with higher investments.

4.2. Analytical Framework

We introduce a simple model as an analytical framework for our study with two main goals in mind: On the one hand to clearly define the concept of trust that we wish to work with and on the other hand to integrate some industry specific observations that we encountered in the course of the qualitative part of our research project into the canonical model. In particular, we want to analyze how competition between suppliers—induced by the OEM—relates to issues of underinvestment due to hold-up problems.

What we observe in reality when a new part is to be developed and procured is that in most cases multiple suppliers are invited to develop a blueprint to match the OEM's specifications. Depending on the development stage, these can range from extremely vague (pre-development up to 3 years in advance of series production) to rather specific (detailed development about 6-9 months prior to series production). A large part of the compensation of the supplier is comprised of the OEM's subsequent business regarding series production or more detailed development—but in these later phases only a subset (often only one firm) of the previously employed suppliers will be awarded the second contract.

4.2.1. Structure and Timing of the Procurement Game

We model this in the following way: Assume that a monopolistic OEM needs the input of a supplier in order to produce output with a value of $v(\theta_i, I_i)$, where the arguments of v designate the intrinsic quality and effort choice of the supplier to be defined exactly below.

First, in $t = 0$, the OEM irreversibly[9] chooses n ex-ante identical suppliers, inviting them to develop and submit a blueprint for the part it wants to procure. For each supplier, it incurs costs k, e.g. administration and coordination efforts, so that the total costs are nk. The suppliers then each independently

[9]Imagine that the interfaces between the part in question and the other parts have been designed by the time the OEM could reconsider, making this option economically impossible.

draw an intrinsic quality parameter θ_i from the known distribution $Q(\theta)$, which is continuously increasing within the domain $[0, 1]$. There are many ways to interpret this parameter - for example it could measure how well the supplier's employees involved in the project are compatible with the OEM's engineers, or how well the research capacities of the two firms complement each other with regard to the period of time in question. Clearly there is a certain level of randomness involved here, even if the firms have cooperated already in the past, and we would argue that in an established industry like the automotive industry, this randomness should be of similar or the same degree over comparable suppliers. After each supplier privately observes the value of its draw, it choses an effort level I_i, given the revelation of θ_i and the n chosen by the OEM. Exerting effort induces the costs $c(I_i)$, with both c' and c'' strictly positive.

Then, in $t = 1$, the suppliers' qualities and investment choices become common knowledge. We limit the subset of suppliers that receive a second contract to a singleton, which is selected in the following fashion. If $n = 1$, then this supplier is awarded the contract, otherwise, the suppliers engage in a second-price auction (or equivalently Bertrand competition) under perfectly symmetric information. Intuitively, this ensures that the supplier contributing the highest net value is awarded the contract.

Finally, in $t = 2$, the supplier with the winning bid and the OEM bargain about how to share the joint surplus from production $v(\theta_i, I_i)$, which is strictly increasing both in θ_i and in I_i. In this bargaining process, we denote the OEM's outside option as g, which is exogenously given. The supplier's outside option is endogenously determined by its effort choice and will be denoted $f(\theta_i, I_i)$.[10] We make the following simplifying assumptions with regard to the relationship of the joint surplus and outside option $f(\theta, I)$ of the suppliers:

A1: $f(\theta, I) < v(\theta, I)$ and $\frac{\partial f(\theta,I)}{\partial I} < \frac{\partial v(\theta,I)}{\partial I}$ for any given (θ, I).

This assumption implies that the surplus generated from the supplier's in-

[10] In an excursion below, we will propose an alternative setting inspired by discussions with industry representatives, in which the outside option is replaced by the OEM covering a share of the suppliers' costs.

vestment within the relationship is always greater than the surplus obtained from that investment outside the relationship. Further, an increase in investment by the supplier leads to a greater creation of value within than outside the relationship.

A2: $v(\theta, I) - f(\theta, I) > v(\theta', I) - f(\theta', I)$ if $\theta > \theta'$.

Higher quality suppliers generate a larger surplus above their outside option than lower quality suppliers for a given level of investment.

A3: $v(\theta, I) > v(\theta', I') \Leftrightarrow v(\theta, I) - f(\theta, I)) > v(\theta', I') - f(\theta', I')$

The higher the total surplus within a relationship, the higher is the efficiency loss suffered if bargaining breaks down. This implies that a supplier with a lower intrinsic quality θ can overcome this disadvantage to become more efficient through higher initial investment I. **A1** through **A3** would be satisfied for an exogenous outside option $\hat{f}$ of the supplier. Given these assumptions, clearly the first best investment level is determined by the first order condition $v_2 = c'$. We will call any level of investment below the first best *underinvestment*.

We distinguish two polar subcases with respect to the bargaining situation at $t = 2$. In subcase a), which we will refer to as the ***de facto*** scenario, the OEM has the entire bargaining power and is in the position to make a take it or leave it offer to the supplier. In subcase b), which we call the ***de jure*** scenario, it is the supplier who is endowed with the superior bargaining position and who makes the take it or leave it offer. One possible interpretation is that the supplier obtains the property rights of the blueprints it has developed at $t = 1$. In the *de jure* scenario the OEM honors these property rights, while in the *de facto* scenario he is able to extract rents from the supplier beyond his formal legal rights without having to fear repercussions.

4.2.2. Results

We solve for a subgame-perfect equilibrium by backward induction for each of the two bargaining settings.

Case a) *de facto* - OEM makes take it or leave it-offer in period $t = 2$

At the bargaining stage in $t = 2$, the OEM offers the supplier a fixed payment T_2, maximizing $v(\theta_i, I_i) - T_2$. He has to take the supplier's participation constraint into account, therefore he must offer at least the outside option such that $T_2 \geq f(\theta_i, I_i)$. This is the only constraint of the problem at this stage, therefore the OEM will choose exactly $T_2 = f(\theta_i, I_i)$.

Now consider the competitive situation among suppliers in $t = 1$. If a supplier is awarded the subsequent contract, we have just seen that he will obtain an offer from the OEM that is equivalent to his outside option, i.e. the OEM extracts the entire rent from cooperation and the supplier is indifferent between winning or losing the bid.[11] Nevertheless, there is a substantial effect of n from the point of view of the OEM. As suppliers are indifferent and there is symmetric information at this point, he is able to select the highest-quality supplier out of n. Mathematically, this is a simple order statistic problem. We know that the maximum order statistic with n draws follows a distribution $Q_1^n = [Q(\theta)]^n$. For non-degenerate Q, this stochastically dominates the original distribution, therefore $\partial E_\theta(n)/\partial n \geq 0$ and increasing the number of suppliers n leads to a higher expected quality θ_1 of the highest quality supplier among the competitors, without affecting their investment incentives.

As a result, each supplier faces the following maximization problem in $t = 0$: $max_I f(\theta_i, I_i) - c(I_i)$, leading to the first order condition $f_2 = c'$. By **A1** the suppliers' investments are more valuable within the relationship than outside it. In this setting therefore the suppliers underinvest. If we denote the supplier's optimal investment level in this case as $I_a(\theta_i)$, the ex ante expected profit of the OEM is:

$$E\Pi_a^{OEM} = E[v(\theta_1, I_a(\theta_1)) - f(\theta_1, I_a(\theta_1))|n] - nk \tag{4.1}$$

where the subscript 1 indicates that it is the expected quality and investment

[11] The winning bidder receives his outside option as the outcome of $t = 2$ bargaining, the losing suppliers receive it because they have to fall back on it.

of the highest quality supplier that is relevant. The OEM benefits from larger n only through the higher expected quality of the best supplier and extracts the entire value added through cooperation. Therefore the optimal n_a^* is chosen at the point where the difference between marginal benefits from expected profit gains due to higher expected quality and administrative costs becomes negative for $n_a^* + 1$.

Case b) *de jure*, supplier makes take it or leave it-offer in period $t = 2$

To clearly demonstrate the effects we are interested in, we first consider the case $n = 1$, in which a single supplier has been chosen originally. Then, in $t = 2$, the supplier offers the OEM his outside option, which we denote by g, according to the same rationale as in the previous case. We make the assumption that the OEM's outside option does not depend upon the effort provided by the chosen supplier, which can easily be justified.[12] Then in period $t = 0$, the supplier's maximization problem is $\max_I v(\theta_i, I_i) - g - c(I_i)$, leading to the first order condition $v_2 = c'$. Obviously, this contract induces first best effort choice by the supplier, who in this case owns the entire project.

The case in which more than one supplier compete amongst each other in $t = 1$ differs substantially from this. Again, first consider the supplier's bargaining problem in $t = 2$ who prevailed in the previous competition with the winning bid, which we call b_1. As the supplier owns the property rights to the blueprints, this bid can be interpreted as the rent granted to the OEM by the winning supplier. How high will the suppliers be willing to bid? The winner's payoff – ignoring the investment costs that are sunk at this stage – will be the total surplus of the project minus his bid, i.e. $v(\theta_i, I_i(\theta_i)) - b_i$, while the losing bidders still obtain their outside option worth $f_j(\theta_j, I_j(\theta_j))$. By **A1**, therefore, the suppliers derive higher rents from winning the competition, as long as the amount they have to pay is smaller than $\hat{b}_i = v(\theta_i, I_i(\theta_i)) - f_i(\theta_i, I_i(\theta_i))$.

[12] The results from contacting another supplier do not directly depend on the effort the first supplier has provided, for example.

At $t = 2$, the suppliers compete à la Bertrand or in a second price auction. For simplicity, we assumed that at this stage, the individual qualities and investment levels are common knowledge.[13] We obtain the well known result for this competition that in equilibrium the most efficient supplier i, with $i = argmax_{i \in N} v(\theta_i, I_i(\theta_i)) - f_i(\theta_i, I_i(\theta_i))$, submits an offer to the OEM, at which the second most efficient supplier is indifferent between matching the bid and her outside option.[14]

To understand the economic mechanisms at play here, it is helpful to consider the OEM's payoffs at given levels of investment and given the choice of suppliers n. If we denote the most efficient supplier as supplier 1 and the second most efficient as supplier 2, the OEM's payoff at given investment levels is simply: $v(\theta_2, I_2(\theta_2)) - f_2(\theta_2, I_2(\theta_2)) - nk$, which is equivalent to the surplus the *second most* efficient supplier would receive if running the project, minus the costs of soliciting the n offers. Ignoring the investment incentives for the moment, two effects are at play: as in case a) with increasing n the expected quality of the second best supplier improves – though it follows the second instead of the maximum order statistic distribution. We consider this a level effect – the OEM profits from the overall expected quality increasing. For the second effect in question, it is useful to first derive the surplus of the most efficient supplier net of his outside option and ignoring the sunk effort costs incurred in the previous period:

$$(4.2) \qquad [v(\theta_1, I_1(\theta_1)) - f(\theta_1, I_1(\theta_1))] - [v(\theta_2, I_2(\theta_2)) - f_2(\theta_2, I_2(\theta_2))]$$

While in case a) the OEM by construction absorbs the entire surplus from cooperation no matter how many suppliers compete, now the most efficient supplier also receives a share, which is determined by the difference in surplus

[13] This assumption is not necessary in this setting, as bidding their real valuations is a weakly dominant strategy in the second price auction.

[14] We obtain the same outcome with a second price auction – here each supplier bids her valuation and the most efficient one pays the second most efficient supplier's bid.

between his and the second best design. As n becomes larger this difference shrinks, so that the OEM receives a greater share of the surplus, which is the second effect we alluded to above.

Up until now, we have taken the investment level of suppliers as given to examine the effects of competition on the bargaining outcome. Now we turn our attention to the investment decision of suppliers in $t = 0$. First let us consider the expected payoffs *given that* the supplier is awarded the production contract later on, conditional on the number of suppliers involved in the competition and the realization of his own quality θ_i. Again denoting the most efficient supplier as supplier 1 and the second most efficient supplier as supplier 2, these are :

$$E[(v(\theta_1, I_1(\theta_1)) - v(\theta_2, I_2(\theta_2)) + f_2(\theta_2, I_2(\theta_2))|n, \theta_1 = \theta_i] - c_i(I_i))$$

For notational simplicity, we abbreviate this expression to $E(S_i|n, \theta_i) - c_i(I_i)$. We can take the surplus generated by supplier 1 out of the expectation operator, which means that his profits *given that he receives the contract* are the value he generates, reduced by his investment costs and the expected surplus generated by the second best supplier. As noted above, the latter increases in the number of suppliers involved for given levels of investment. But supplier i does not receive this payoff with certainty, but instead he anticipates that he will be awarded the subsequent contract only if he is the most efficient supplier. As at this point the intrinsic quality of the other suppliers is unknown, the supplier expects this event to take place with probability $p[v(\theta_i, I_i) - f(\theta_i, I_i) \geq \sup_j v(\theta_j, I_j) - f(\theta_j, I_j)]$ with $i \neq j$. Applying **A3**, this can be reduced to the shorter expression $p[v(\theta_i, I_i) \geq \sup_j v(\theta_j, I_j)]$.

Now, spelling out the maximization problem of supplier i in $t = 0$, we get:

$$\max_I p(\theta_i, I_i, n)E(S_i|n, \theta_i) + [1 - p(\theta_i, I_i, n)]f_i(\theta_i, I_i) - c(I_i) \tag{4.3}$$

This leads to the first order condition with respect to optimal investment:

(4.4)

$$c'(I_i) = \frac{\partial p(\theta_i, I_i, n)}{\partial I_i}\left[v(\theta_i, I_i(\theta_i)) - \max_{j\neq i} v(\theta_j, I_j(\theta_j), n) - f_i(\theta_i, I_i)\right] + p(\theta_i, I_i, n)\frac{\partial(v(\theta_i, I_i) - f_i(\theta_i, I_i))}{\partial I_i} + \frac{\partial f_i(\theta_i, I_i)}{\partial I_i}$$

We observe a number of countervailing effects. On the one hand, investment incentives are somewhat diluted, as the supplier only profits from the added value with some probability. On the other hand, he can increase the probability that the contract will be awarded to him by choosing higher levels of investment. Note that irrespective of the number of suppliers n, the investment incentives are strictly higher in this case than in case a) – intuitively, in the *de facto* scenario, the suppliers *never* participate in the surplus they generate within the relationship, therefore even the incentives diluted by competition induce a higher effort level. Only as n gets very large, the two effort levels converge at $f_2 = c'$.

Again denoting a supplier's investment choice in this subcase analogously to above as $I_b(\theta_i)$ and the most and second most efficient suppliers as supplier 1 and supplier 2, respectively, the expected profit of the OEM in this case is:

$$E\Pi_b^{OEM} = E[v(\theta_2, I_b(\theta_2, n)) - f_2(\theta_2, I_b(\theta_2, n))|n] - nk \quad (4.5)$$

Comparing equations (1) and (5) may lead to the impression that the OEM's profits are necessarily higher in the *de facto* case, as one is comparing the maximum and the second order statistic. But unlike case a), in which the supplier's investment level was not affected by the choice of n, here the supplier's investment choice does depend on the level of competition and is, at the same time, strictly larger than above. As a result, we cannot in general rank the two profit levels.

We are able to say a bit more about the optimal number of suppliers in the two scenarios. In a), increasing the number only has the effect that the expected quality of the best supplier increases. In b), we have three effects: 1) The OEM expects a higher quality, 2) he receives a greater share of the surplus as n increases, and 3) the suppliers invest less into quality. It is extremely simple to show that effects 1) and 2) dominate the quality effect in scenario a), i.e. would lead to a higher choice of n, were it not for the investment effect. Then, we really are only comparing the maximum with the second order statistic – we know that as $n \to \infty$, the second order statistic approaches the maximum order statistic from below. As both are monotone increasing, the slope of the second order statistic must be strictly greater. As a result, the OEM will choose a higher n in case b) than in case a) unless the investment-deterring effect 3) is "too large" in relation to 1) and 2). To enhance the intuition for this result, consider the following simple example: Let supplier quality be uniformly distributed so that $Q(\theta) = \theta$. The suppliers' investment cost function is determined by $c(I) = \frac{I^2}{2}$. The surplus derived from cooperation is $v(\theta, I) = \theta I$ and each supplier's outside option is worth $f(\theta, I) = av(\theta, I)$ with $a \in (0, 1)$. Finally, the OEM's outside option is $g = 0$. Here we are able to derive closed-form solutions which we summarize in the following table:

	$I(\theta)$	$E(v)$	$E\Pi^{OEM}$
a) $n = 1$	$a\theta$	$\frac{a}{3}$	$\frac{a(1-a)}{3} - k$
a) $n > 1$	$a\theta$	$\frac{an}{n+2}$	$\frac{a(1-a)n}{n+2} - nk$
b) $n = 1$	θ	$\frac{1}{3}$	$-k$
b) $n > 1$	$(1-a)\theta^n + a\theta$	$\frac{(1-a)n}{2n+1} + \frac{an}{n+2}$	$\frac{a(1-a)n(n-1)}{(n+1)(n+2)} + \frac{(n-1)(1-a)^2}{2(2n+1)} - nk$

Table 4.1.: Investment decisions, expected surplus and expected profits of the OEM in $t = 0$
$\big($for $Q(\theta) = \theta$, $c(I) = \frac{I^2}{2}$, $v(\theta, I) = \theta I$, $f(\theta, I) = av(\theta, I)\big)$.

In the *de jure* case, the suppliers' investment is higher by $(1 - a)\theta^n$ than in the case in which the OEM extracts all surplus ex post. With increasing n, these investment incentives decrease, but due to the functional form, for $\theta \to 1$, the pace of this process goes towards zero. As a result, a "'wedge" between the expected value v of the two strategies remains even as the number of suppliers grows very large.

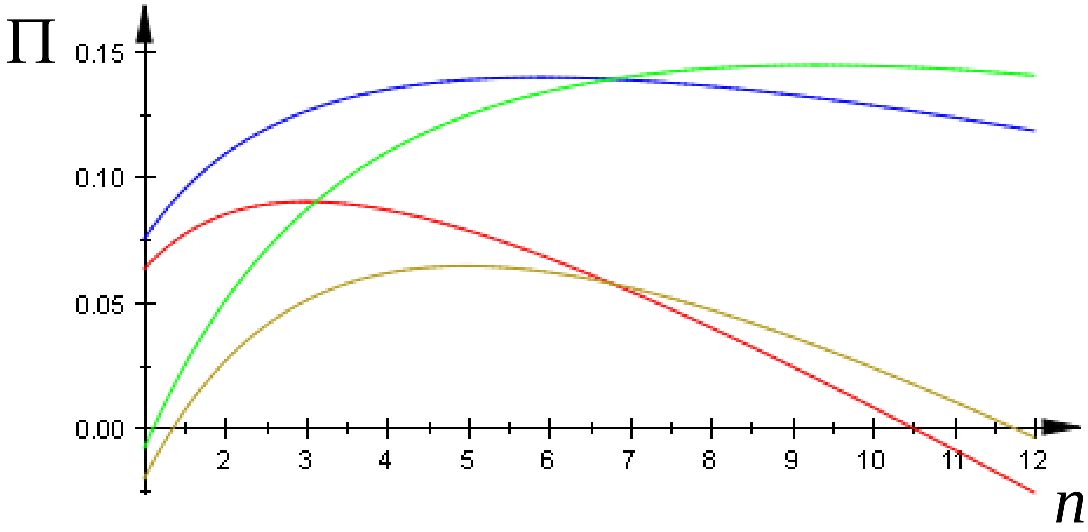

Figure 4.1.: OEM's expected profits depending on n
(Blue: Π_a (low k); red: Π_a (high k); green: Π_b (low k); brown: Π_b (high k)).

What we care about at this point, though, is the optimal choice of n. For this, we can compare the expected profits in the third column. First, compare only the first terms of $E\Pi_a$ and $E\Pi_b$. The term in the *de jure* case starts out at 0 and approaches the level of the first term in the *de facto* case as $n \to \infty$. For the given functional form, this isolates the effect of the second order statistic approaching the maximum order statistic at given investment levels alluded to above. The second term of $E\Pi_b$ is also strictly increasing in n, which reflects the higher share of the (larger) total surplus being absorbed by the OEM. Therefore $\frac{\partial E\Pi_b}{\partial n} > \frac{\partial E\Pi_a}{\partial n}$ and the OEM's optimal choice of n is larger in scenario b) than in scenario a). **Figure 4.1** depicts the OEM's profits in these two scenarios and gives an example for which they are larger in each of the two settings, respectively, depending only on the relative size of administrative costs k. Note that while it is not clear which strategy is more profitable in general, $n_a^* < n_b^*$ holds.

4.2.3. A Simple Notion of Trust

Trust and Underinvestment

Up until this point, trust does not explicitly play a role in our model. We attempt to integrate the concept in the following way. Assume that prior to their investment choice, it is not certain which subcase, a) or b), will be played later on in the game. Instead, with probability λ it is the supplier who makes the take it or leave it offer in period 2, while with probability $1 - \lambda$ it is the OEM who makes the offer. Correspondingly, the supplier is able to generate a rent above his outside option with probability λ.

How does this relate to the reality in the industry? In our in-depth interviews, the clear picture emerged that OEMs are at great liberty in designing and enforcing contractual details in relationships with suppliers, almost irrespective of the size or market power of their counterparts.[15] Ben-Shahar and White (2006) report equivalent or even more pronounced findings for the North American automobile industry. At first glance therefore it would appear as if only the a)-subcase (allotting the entire bargaining power to the OEM) described above are relevant with respect to reality. This confrontational setting, in which the relationship is defined mainly through the pure holdup-problem, is often subsumed under the term of an "American" procurement strategy. Yet the global success of Japanese carmakers beginning in the 80s has prompted much interest in alternative ways of supply-chain management, perhaps most famously incarnated in the MIT's International Motor Vehicle Program. As a result, researchers started to stress the importance of cooperative and mutually beneficial relationships between OEMs and suppliers in the industry in achieving the goals of lean production. In this context, Taylor and Wiggins (1997) show that granting the suppliers positive economic rents can be a substitute for control (i.e. monitoring quality levels and eventual punishment).[16] Traditionally, cooperative relationships are prevalent in the Japanese automotive

[15] See Müller et al. (2008) for details.
[16] See also Aghion et al. (2002) for qualitatively similar results in a more general setting.

industry and were common practice in Germany at least up until the mid-90s.

For a "traditional" micro-economic perspective, consider the entire game from the previous sub-sections to be only the *stage game* of an infinitely repeated Markov-game. λ then is a state variable that depends on the (in our reduced-form unobserved) previous history of the game. At the point in time in which the stage game that we observe is played, λ is given and resembles the subjective probability with which subcase b) will be played in this period, i.e. the probability with which the OEM respects the supplier's property rights. This is precisely our definition of trust in the following: Higher values of λ denote a higher subjective probability with which the OEM will grant the supplier the extraction of a share of the rent generated within the relationship – in other words, the higher λ, the more the supplier *trusts* in the OEM's willingness not to exploit his superior market power.

We refrain from spelling out the repeated game for the following reasons: First, the data on which we base our accompanying empirical analysis are cross-sectional without a panel-dimension. Therefore, what we observe can be considered exactly one period of the stage game, in which the current relationship-specific levels of trust are exogenously given, and this allows us to apply the model more directly. Second, in this setup of the model, what matters is the supplier's level of trust toward the OEM. We believe that the main driving force in these relationships is the hold-up problem and therefore our objective is to see whether and how this can be mitigated by the belief of suppliers that abuse is less (or more) likely. Finally, looking at this reduced form allows us to remain agnostic about the reasons for the OEM wanting to grant the supplier this rent.[17] To state just two examples, as in Taylor and Wiggins (1997) there could be a tradeoff between rent extraction and control costs, or relatedly, it may be worth while to make the supplier fear the consequences of the relationship being terminated as in Akerlof and Yellen (1990).

Our model allows us to state our central prediction very straightforwardly:

[17] For the game to be on the equilibrium path, the subjective and the objective probability of subcase b) being played would have to be the same.

Hypothesis 1: *Higher levels of supplier trust in the OEM are associated with higher relationship specific investments by suppliers.*

As shown above, this holds for any given procurement strategy chosen by the OEM – but even this reduced form model conveys some of the complexity of the procurement decision, which also depends on the level of quality uncertainty regarding the product. Clearly, all else given, higher levels of quality uncertainty will induce procurers to induce competition between more suppliers.

Trust and the Optimal Induced Level of Competition

While we get a clear cut and very intuitive result with respect to the relationship between trust and upstream investment, the analysis with regard to the level of competition induced by the OEM is somewhat more complicated – and surprising. As described above, a large part of the relational contracting literature, see e.g. recently Calzolari and Spagnolo (2009), argues that competitive arms-length agreements on the one hand and relational contracts on the other will be used in mutually exclusive settings, depending, for example, on the enforceability of complex clauses. Supported by these findings one might expect in a naïve first approach that supplier-OEM relationships governed by trust, i.e. with relatively high λ in our model, should be associated with less induced competition, i.e. lower n. Our analysis above has shown that generically the opposite is going to be the case, as λ simply generates a convex combination between scenarios a) and b) and in the latter, a higher level of competition is optimally induced. **Figure 4.2** displays how the optimal n increases for the simple example stated above as λ increases from 0 to 1.

The intuition behind this result is as simple as it is striking. In the *de facto* setting, the OEM extracts the supplier's entire rent through his superior bargaining position, exploiting the existing hold-up situation. In the *de jure* setting, he forgoes this possibility – inducing competition, i.e. a choice of a higher n is his alternative mechanism to extract rent from the cooperation. In this sense, trust and competition are *complementary* in this setting. We can

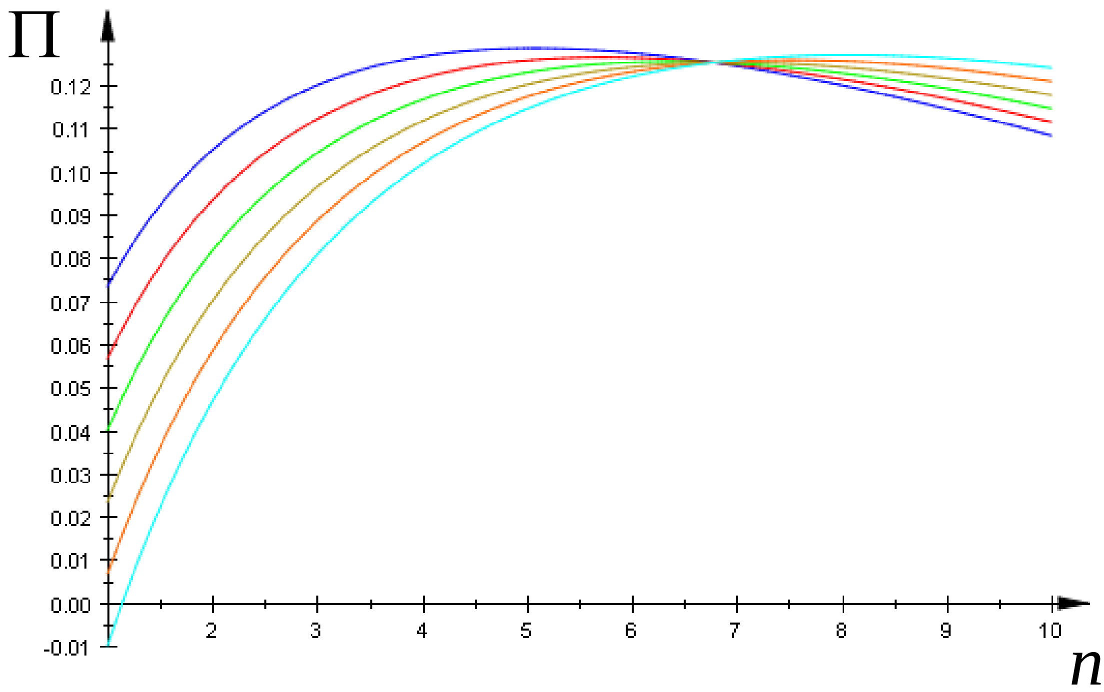

Figure 4.2.: OEM's expected profits depending on n for different values of λ (Blue: 0.0; red: 0.2; green: 0.4; brown: 0.6; orange: 0.8, cyan: 1.0).

directly derive the second hypothesis from the formal model:

Hypothesis 2: *Higher levels of supplier trust in the OEM are associated with more intense supplier competition in the procurement process.*

While this results rather directly and clear cut from the formal model, it will prove to be very difficult to distinguish our concept of trust from, e.g., the concept of relational contracting or a "naïve" understanding of the term. In each of these cases, the predictions would be the opposite of our second hypothesis, and therefore we approach the task at hand as one of competing hypotheses and will let our data decide.

4.3. Empirical Analysis

4.3.1. Source of Data

Our data results from an online questionnaire study that was carried out for the German automotive industry association[18] (VDA) from Fall 2007 until Spring 2008. The questionnaire was designed on the basis of the results of a case study performed in Spring 2007, in the course of which interviews with high ranking executives in the automobile industry were conducted.[19] We obtained a unique view of the relationships between original equipment manufacturers and their tier 1-suppliers, with a twofold approach: First, each of 13 participating suppliers was asked to evaluate their relationship with each of up to 11 OEMs active in the German market for different representative products in their portfolio in clinical detail – more than 300 questions were asked covering all central functions within the firms. In addition, the participating OEMs were asked to evaluate their sourcing relationships in general – i.e. not specifically for individual suppliers – for each of the four different product classes according to the established industry classification:

- **Commodities:** physically small and technologically unsophisticated (e.g. shock absorbers)
- **High-tech components:** physically small but technologically sophisticated (e.g. electronic brake component), in the following referred to simply as component.
- **Modules:** physically large but technologically unsophisticated (e.g front end)
- **Systems:** physically large and technologically sophisticated (e.g. break system)

[18]Verband Deutscher Automobilunternehmen e.V.
[19]For the qualitative results of this case study, see Müller et al. (2008).

As OEMs answered a set of questions almost identical to the supplier questionnaire, we are thereby able to compare their view of their general policies with the suppliers' view. In total, more than 1,500 questionnaires were filled in by competent engineers, procurement- and sales officers with the following methodology. A participant first would have to indicate his function within the company out of the following[20]:

- Pre-development: "Basic" technological research, not model-specific.
- Vehicle Development: Car-model specific (technology adaptation).
- Series Production
- Quality Control
- Sales
- Logistics
- Aftermarket

Then she would choose a product for which she had the necessary know-how as well as the customers she worked with, the latter from a list. For each product and customer, she would then answer a set of questions suited to her function within the company.

One observation in our data is composed of the answers of the *entire* supplier questionnaire for a given product and a given customer. This exceeds the extent of any given function-questionnaire listed above – therefore, in order to obtain as complete observations as possible, we merge the answers from a given supplier, product and customer over all functions to cover all aspects of the relationship. What each observation therefore describes, is one view (potentially of many people) of the relationship between the supplier and a

[20]For a detailed description of the individual functions and the automobile development and production process, we refer to Müller et al. (2008)

given OEM for one product. We merge this with the results from the OEM-questionnaire in order to be able to control for the non supplier-specific behavior of OEMs.

On paper, we have 792 observations, but for two reasons these are not necessarily complete: First, not each function within a company filled out a questionnaire for each product studied. In this case, whole sections of the questionnaire are missing for the observation covering the given product. Second, participants could skip individual questions and made ample use of this option. Therefore the numbers of observations over the individual questions differ substantially, as seen in the descriptive data below.

4.3.2. Descriptive Statistics

The underlying questionnaire sought to depict complex relationships in hitherto unmatched detail. In the following subsections, therefore we will exert effort – perhaps more than usual – to introduce the variables of the study and shine some light on the basic forces and tensions that are at play between manufacturers and suppliers.

Participating Companies and Characteristics of Parts

On the OEM side, 10 of the largest players in the German market participated actively in the survey, 7 producers of passenger cars and 3 truck makers. On the other hand, 13 suppliers active in the German market provided their input on 11 manufacturers - i.e. the 10 participating plus one further car manufacturer. The supplier sample is very strongly biased towards large participants, with average 2007 revenues of 9.4 billion Euro (standard deviation (std) 12.4) and even the smallest participant posting revenues above 700 million Euro. This is emphasized by the self-reported European market shares for the product in our sample: For 161 observations this was provided on a five point scale with an average of 3.76 (std .90), which translates into a share of more than 25%. Not surprisingly, the correlation of market share with intensity of supplier competition – also on a 5-point scale – is negative with a value of -0.20 (significant at 5%-level). Further, we observe a negative correlation of supplier size (measured by the 2007-revenues in billion Euros) and the intensity of supplier competition (-.144, p-Value: .072) and a positive correlation of supplier size and market share in the observed part (.124, p-Value: .083).

This could raise the worry that the larger suppliers may be able to exert monopoly power over OEMs for some of the parts we study – to counter this, we made sure that there are at least two suppliers active in Germany for each part in our sample. Nevertheless, we will have to try to control for relative market power in our regressions, as it may clearly affect bargaining strength and the OEM's outside option.

Apart from their type, products are further specified by the R&D-share of total costs as well as the assessment by the supplier how important the degree of innovation is for the particular part. Both were measured on a 5-point scale – the importance ranging from 1 – very little – to 5 – very high, while the cost-shares were provided in 2% increments, therefore ranging from $< 2\%$ to $> 8\%$. As one would expect, the answers to the questions are strongly correlated (0.27, significant at 0.1%-level). More interestingly, though, they allow us to revisit the merits of the underlying type-classification. The following **Table 4.2** displays the descriptive statistics for these questions by underlying product type:

Performing pairwise t-tests shows that the means for both variables are significantly lower for commodities and modules than for systems and components, while among these two groups the hypothesis of equal means cannot be rejected, which is exactly in line with the industry specification discussed above. This allows us to introduce an additional dummy measure of the innovativeness of a part (dummy_soph) which takes the value 1 if the type is a system or component. Further, to account for potential price differences due to the sheer size of a part, we introduce dummy_big which takes the value 1 for systems and modules.

Variable	Mean (Std. Dev.)	Min	Max	Obs
Systems				
Cost Share R&D	4.19 (1.43)	1	5	37
Importance R&D	4.16 (.74)	1	5	50
Modules				
Cost Share R&D	2.06 (1.76)	1	5	18
Importance R&D	2.89 (.76)	1	4	18
Components				
Cost Share R&D	3.48 (1.40)	1	5	21
Importance R&D	3.42 (.87)	2	5	45
Commodities				
Cost Share R&D	2.35 (1.35)	1	5	91
Importance R&D	2.91 (.69)	2	4	93

Table 4.2.: Importance of Innovation and Cost share R&D by product type.

Characteristics of the Relationship

In this section, we pursue two different goals. While introducing the variables describing the relationships we also attempt to shed additional light on the differences pertaining to product-type. There is a further dimension, in which we can exploit existing variation: As noted above, we gather information on the relationship between OEMs and suppliers over the entire car-model life-cycle. In the following, the distinction between three of these phases is especially important: pre-development, development and series production. The last of these phases, series production, is the least complicated case – suppliers work with existing blueprints and completely designed (or existing) tools to produce given quantities of the part in question. The product and services can clearly be specified through contracts without much room for misunderstanding, for example specifying acceptable failure rates and delivery conditions in detail. The (model-specific) development phase is in many ways less clear cut. While the general requirements that a part has to meet are defined by its function within the automobile (a brake has a relatively specific function and place, given the projected weight and top-speed of the model in planning), a plethora of other parts with which it has interfaces are being designed in parallel. Blue-prints for the part do not yet exist at the beginning of the design phase. Clearly, the objectives cannot be drawn up precisely ex ante in contracts, but are subject to a continuous cooperative process. Finally, these uncertainties become overwhelming when considering the (not model-specific) pre-development phase: Here for example, the supplier is researching brake-technologies without knowing how fast or heavy the model in which it will be used is going to be. In general, more fundamental research is involved here – and, as should be clear from the nature of the endeavor, it is even harder to write specific and precise enforceable contracts regarding the outcomes. As our respondents were involved in the different stages of the product life-cycle, in a way like Brown et al. (2004) this allows us to exogenously change the level of external enforceability while keeping product and relationship characteristics constant for a number of questions.

As a case in point, we requested suppliers to evaluate the OEM's supplier choice criteria on a six-point scale from 1–no relevance to 6–very important, for each of the phases in the product life-cycle. From the discussion in the literature we would expect "relational" choice criteria to grow relatively less important as opposed to "hard" criteria – such as price – as one progresses from pre-development onward. Our empirical results strongly support this hypothesis. The importance of price strictly increases from an average of 5.10 (see Table B.1), to 5.37 (see Table B.2) to 5.70 (see Table B.3) (t-tests for difference of mean are each significant at the 0.1% level). On the other hand, the importance of trust is respectively 4.89 (see Table B.1), 4.90 (see Table B.2) and 4.73 (see Table B.3) for pre-development, development and series production. Therefore trust is only significantly less important when choosing a series supplier ($p<0.1\%$), while there is no difference between pre-development and development. But with this kind of question, it is easier to interpret relative magnitude of answers: When we look at the differences in differences between the importance of the choice criteria, there is a monotone relationship, with price becoming relatively more important for each step ($p<5\%$ for pre-development to development and $p<0.1\%$ from development to series production).

[**Table B.1, B.2** and **B.3** about here]

Next, we suggest two proxies for the value of the OEM's outside option (denoted g in the theoretical framework). The first is the share in the volume of the part provided by the supplier to the OEM (measured on a 5-point scale where each point resembles a 20% difference, with 1–<20 and 5–>80%). Presumably, it is more difficult to shift a larger share of production away from one supplier to another than a smaller one, therefore g should be negatively related to this measure. Looking at how this measure behaves for the different product types shows that the share provided has statistically non-distinguishable means for modules, components and commodities (the sub-sample means are, respectively, 3.08, 3.26 and 3.23, i.e. at the upper end of the range 40-60%), while for the systems, this value is significantly higher at 4.07 (or at the upper end of the range 60-80%). The reliance on an individual supplier therefore is significantly stronger in the case of systems than for the other types of part.

As the second potential proxy, the respondents were asked to assess how often the OEM chooses to produce a given part himself on a 6-point scale from 1–never to 6–very frequently, with 4–about 50% of cases as a further anchor. This also allows us to create a dummy variable which takes the value 1 whenever the answer is different from *never*. The ability to produce a part himself is perhaps the most intuitive outside option – and it is one of which OEMs are making ample use, as the recent wave of "in-sourcing" demonstrates. In day-to-day business, there can be different reasons for this, most commonly capacity utilization smoothing[21] or worries about suppliers' ability to provide a part as agreed. The comparison of sub-sample means shows that systems and components, i.e. the technologically sophisticated parts, are significantly less likely to be also produced by the OEM himself than commodities and modules, with the latter being the most likely to be in-sourced. This may be due to the fact that physically larger parts are more costly to stockpile and utilize

[21]Labor laws make short-term adjustments to the workforce all but impossible.

more capacity. Clearly, higher levels of this variable should be associated with a higher outside option of manufacturers, i.e. a higher g.

Finally, we asked the respondents to evaluate the level of specificity of the contractual requirements at different development stages, as well as the “degree of freedom” in relationship in order to be able to generate some insights into the interactions between trust and contract specificity.

Procurement Decisions by the OEM

We have two sets of variables that measure the OEM’s procurement decisions at different points in the product life-cycle, one qualitative and one quantitative.

For the qualitative measure, we asked the respondents to evaluate how often different procurement strategies have been employed by the OEM for each of the different stages. This may appear slightly paradox, as for each part a manufacturer should apply one strategy, but parts are procured anew for each new series of a given model, i.e. there is a new procurement process every 1.5 to 2 years and clearly different strategies could be used at different points of time in the past. For pre-development, the options were *preselection of a specific supplier* and *procurement among a limited number of suppliers*, each on a 6-point scale from 1–never to 6–very frequently. For development and series production, a further option was added, *open procurement*, which plays no role in pre-development. Even the purely descriptive results offer some interesting insights. For pre-development, OEMs are actually significantly more likely to contract with specific suppliers (mean 4.43), than to go through a limited competitive procurement process (mean 3.95, t-test for difference of means significant at 1% level; see **Table B.4**). In contrast to this, preselection of suppliers is significantly less likely both for development (mean 3.06) and series production (2.98), see **Tables B.5** and **B.6**, respectively. On the other hand, for development OEMs are significantly more likely to procure among a limited number of suppliers (mean 5.18, see **Table B.5**), therefore there is a clear shift to more market-based interactions from pre-development to development. The same kind of shift takes place again from development to series production, where procurement among a limited number of suppliers grows less important (mean 4.55, see **Table B.6**), but there is a significant increase in the use of open procurement (2.44 instead of 1.97, see **Tables B.5** and **B.6**). Clearly, the picture that has begun to emerge above, i.e. a shift to more arms-length interactions as the product reaches the development and series production phase is supported by these data. We believe this to be driven mainly by increasing contractibility, when viewed together with the results of the quantitative measure.

[**Table B.4, B.5** and **B.6** about here]

For this, we asked how many suppliers provided the given service or produce the part in parallel, differentiated for additional phases within each of the (by now) familiar stages. The development stage was subdivided into the phases product planning, product specification, concept development and detailed development (starting from the earliest). For series production, we asked for the number of suppliers at series start, after 1-2 years and after more than 2 years. The results from this appear to be somewhat counter-intuitive. For pre-development there are on average more than two (2.16, Table B.4) suppliers competing in parallel. This number stays about constant in the first stages of development, before it significantly *decreases* for the last development phase down to 1.51 (see Table B.5). It reaches its nadir at the beginning of series production with 1.20, before it increases again to 1.59 two years into production (see Table B.6). How does this mesh with our prior results? The previous questions only aimed at the choice procedure, instead of at how many suppliers are selected. For pre-development, due to the lack of specificity concerning the objectives, open procurement is not feasible – precisely for this reason, there is the greatest uncertainty regarding the outcome of the process. The way that OEMs deal with this – also suggested by our model – is to have multiple suppliers work on the designs. As seen above, these are frequently hand-picked. On the basis of the most promising approach the OEM then enters into the development process. There is a strong incentive for suppliers for their preliminary design to be chosen, as the contractual reimbursement for pre-development work is on average below 60% of the actual costs, whether or not the company is awarded a subsequent development contract. An analogous process is repeated again for the development process, which results in a specific blueprint. With this blueprint, the quality uncertainty is practically eliminated, given that suppliers are generally certified through stringent quality assurance processes, therefore this component is eliminated from the decision problem. In production, fewer suppliers with higher volumes promise the highest economies of scales and the steepest learning curves, therefore the number of suppliers drops significantly at production start. Once these effects have been realized, the OEM can start to bring additional suppliers in.

We believe that this background is extremely valuable and should be born in mind for the following analysis of which role trust plays in the interactions set out above.

4.3.3. Measures of Trust: Who Trusts Whom - and Why?

Trust is a sensitive concept which has proved to some degree elusive to attempts at explanation by economists. While existing studies have employed either experimental/behavioral evidence or subjects' answers to variations on the question "To which degree can other people be trusted?", our data has the huge advantage that it is relationship-specific: We ask representatives of company A about their stance and misgivings toward company B with regard to the interactions concerning a specific product. Clearly, there are drawbacks to this approach as well that need to be addressed. We devote the following three subsections to determine how robust the individual measures are and whether and how they can be applied. First we introduce the questions that we believe to be related to the concept "trust". Then we use a method well-established in sociology, exploratory factor analysis, to try to shed additional light on the dimensionality of the construct we are observing – i.e., is there only one kind of "trust", or do the questions we observe really depict a construct composed of various different "factors". Finally, we try to pry the black box from **Section 4.2.3** open (if only a slit), by taking a glance at which past outcomes and behavior affect the suppliers' evaluation of trust in the relationship.

How Do We Measure Trust?

We attempt to capture trust in the various supplier-manufacturer relationships in two, not necessarily mutually exclusive ways. One approach is relatively direct, asking to evaluate mutual trust or inquiring which role trust has played for important decisions. The second approach, which we call indirect, is to look at reported behavior for which trust can be considered a prerequisite.

An important context for trust in these relationships is the area of intellectual property. Especially basic, non model-specific research resembles an important share of suppliers' capital and embodies their ability to differentiate themselves – this ability all but disappears, for example, if an OEM were to take a supplier's blueprints for a part and make them accessible to competitors. Much of this know-how is involved in the earliest stages, the pre-development of products, where suppliers showcase their advanced know-how. In our interviews in the preparation of the study, there was a mention of the practice to deny the most advanced technology to OEMs who were expected not to treat it with the necessary care, i.e. who were not trusted.

Therefore as a first **direct measure** of trust, we inquired after the *importance of the trust relationship with the OEM in a firm's decision to initiate a pre-development project* on a six-point scale ranging from 1–no relevance to 6–very high relevance (from now on **Trust 1**). To be able to relate this to other criteria, we asked the same questions for the importance of the factors *sales potential*, *product positioning* and *long-term cooperation*, so that we can use both the absolute value of the answer as well as the relative rank as measures of trust.

As the second direct measure, we asked the question: "*How do you evaluate mutual trust between OEM and supplier with respect to honoring each other's intellectual property rights (IPR)?*" on a five-point scale ranging from 1–very little to 5–very high (from now on **Trust 2**). While the first question only involves the level of trust of the supplier towards the OEM, the second question is phrased to cover bi-directional trust. Clearly, a disadvantage to this second question is that the supplier must also give an estimate of the other party's assessment.

We already encountered the third direct measure of trust above. For each phase of the product life cycle, the suppliers were asked for their view of the OEM's choice criteria for choosing his supplier (pre-development: **Trust 3**, development: **Trust 4**, series production: **Trust 5**). Again, we have both the absolute value of the importance of trust as well as the relative rank compared to cost, personal contact, duration of cooperation and certification. This is the *supplier's* assessment of the *OEM's* preferences, only, so we clearly need to evaluate the reliability of this measure. Second, analogous to Sapienza et al. (2007), we will have to check whether respondents perhaps related their own level of trust in the OEM in these questions, instead.

Our potential **indirect measures** of trust are associated with behavior that is related to IPR protection and to the secrecy of the cost-structure of the supplier.[22] Suppliers state both *how often they provide original research data to the OEM* on a five-point scale (1–very rarely to 5–very frequently) as well in a separate question *how often the OEM provides access to his original research data* on the same scale. Clearly, both the levels and the difference between the two values may be of interest. Further we inquire on the same scale *how often the supplier's costs are made transparent to the OEM*. An interesting issue with these measures which we will have to attempt to disentangle is whether suppliers are forced into revealing these data due to the OEM's superior market power or whether this is truly a result of trust. In order to determine this, the relation between the frequencies with which a supplier and the OEM reveal original research data will be of interest and we introduce the difference between

[22] The suppliers' costs are an extremely contentious issue in negotiations. Cost-cutting manufacturers (have to) accept that a supplier producing below cost will have to go out of business sooner rather than later. Therefore, they traditionally try to negotiate prices that are as close to the costs as possible and begrudge the suppliers any positive margin that they obtain. For an excellent and comprehensive discussion, see the classical Womack et al. (1991).

the two as an additional variable.

As a first step towards better understanding these measures, **Table B.7** displays pairwise correlations between each of them. As one would have expected, there is a significant positive correlation between all of the direct measures of trust. For the indirect measures, the picture is more interesting: The frequency with which the supplier makes his costs transparent is *negatively* correlated to two of the direct measures – the importance of trust (from the OEM's view) in choosing his development and series supplier. These measures reflect the attitude of OEMs in the selection and negotiation process of development and series suppliers – it appears plausible that it is the insistence of the OEM, therefore, and not necessarily trust that causes suppliers to bare their costs more frequently, which makes this measure non-satisfactory.

[**Table B.7** about here]

The supplier's provision of original know-how, on the other hand, is not correlated with any of the direct trust measures, while there is a relatively strong significant positive correlation (.443) with the provision of know-how by the OEM. The latter is also positively correlated with the mutual trust regarding the treatment of IPR. What would we expect a "trusting" as opposed to a "forced" relationship to look like? If the OEM forces the supplier to reveal intellectual property secrets, this should negatively affect the level of mutual trust with regard to IPR. Further, we would expect that – in these kinds of relationships – the OEM provides relatively little intellectual property into the relationship himself. Finally, the relative market and bargaining power may play a role in this kind of relationship. To determine, whether this effect truly exists in the data, we regress the difference of IP-secrets provided onto the level of mutual trust with regard to IPR, a dummy whether the product is technologically sophisticated and the supplier revenues as a proxy for relative market power. The results of the OLS-regression are provided in **Table 4.3** below.[23]

[23] We also performed ordered logit regressions, which are more suited to the structure of the data. The results

Variable	Coefficient	(Std. Err.)
Trust 2	**-0.233****	(0.113)
dummy_soph	0.249	(0.201)
Supplier Revenue	**-0.015****	(0.007)
const.	2.009	(0.413)

Table 4.3.: (OLS): Difference in frequency of revealing original research data
(N = 129, * significant at 10%, ** significant at 5%, *** significant at 1%)

The regression results show a significant negative association of the difference in IP-provision and mutual trust as well as supplier revenue, which, as we showed above, can be used as a proxy for relative market power.[24] Therefore, the lower the relative market power of a supplier, the more likely it is that he provides more original research than his opposite, which we take as an indication that in these asymmetric setting, enforced revelation (which is negatively associated with trust) does take place. In the simple pairwise correlations the overall effect of supplier provision of IP on our measures of trust is neutral (not significantly different from zero). Therefore there must be information in this measure that countervails the effect of the on average increasing difference in provision of IP. Intuitively, one could imagine there to be three coexisting IP-regimes: One symmetric one characterized by distrust – here, both parties provide little or no research findings to each-other. One asymmetric one characterized by force, the existence of which is suggested by the regressions above. And finally, one symmetric one characterized by trust – here, both parties provide research to each other relatively often and in similar amounts. We try to use the following measure to be able to account for the differences in the three regimes: First we create a dummy that takes the value 1 if the difference between the provided IP is not too large, i.e. no larger than 1.[25] Next, we interact

are qualitatively identical (signs and p-values), we report the OLS regression for the easier interpretability of the coefficients.

[24]The corresponding regression with mutual trust as a dependent and the difference as an independent variable shows a significant negative effect of the difference on trust.

[25]Taking this as a not too large difference is somewhat arbitrary. There are two reasons why we find it sensible. People tend to overestimate their own contribution compared to others. Further, the value of 1 leads to the highest correlation of the final measure with our direct trust measures.

this dummy with the frequency of IP-provision by the supplier.[26] The resulting measure has a significant, positive correlation with the reported mutual trust with regard to IPR (.163, p-value .05).

As a result of these considerations, we are left with the 5 "direct" trust measures. Our knowledge of the relationship between the observed measures up until this point is based on pairwise correlations alone. Factor analysis is a method designed to make better use of these "within" correlations between a set of variables in order to extend what can be learned from them. Using a latent variable approach, it maximizes and records the share of variation in the observed variables that can be explained by one unobserved factor (or more), while reproducing the correlations between variables.[27] The method has been criticized in the past for producing results that are not unique, but we find it perfectly suited to produce a kind of "upper bound" in our exploratory setting, i.e. to explain how much of the variation in our measures can at most be explained through the unobserved underlying factor, which we assume to be (at least associated with) trust. One remaining difficulty that we face is that we only have 59 observations in which all 5 variables are included, but even this low number of observations can be sufficient in a 1-factor, 5-variable model as MacCallum et al. (1999) argue, and we perform a number of tests for robustness.

Table 4.4 displays the factor loadings and uniqueness of the individual variables using the principal-factor method and limiting the admissible number of factors to 1. The resulting pattern is robust to using the maximum-likelihood estimation approach, to allowing a second and third underlying factor and to recursively eliminating individual factors (thereby obtaining significantly more observations). In all specifications, the uniqueness for the variables that measure the importance of trust in procurement negotiations at different stages (trust 3-5) is close to or below the level of .5, which is seen as the relevant threshold in the literature. Among these, the explanatory power regarding pre-

[26]This interaction term has the lowest value for the "force" regime (0), low values for symmetric mistrust regimes and the highest values for symmetric trust regimes.

[27]For an introduction to Factor Analysis, we refer to Harman (1976).

variable	Factor loading	Uniqueness
(1) Trust1	.594	.648
(2) Trust2	.473	.776
(3) Trust3	.679	**.539**
(4) Trust4	.844	**.288**
(5) Trust5	.771	**.406**

Table 4.4.: Factor Analysis Results
Factor loadings and uniqueness reported, principal-factor method (N = 59).

development negotiation (trust 3) seems to be smallest. The general rules of thumb would suggest to remove all variables except for (trust 3-5) from the model.

For us, this entails the following result: It appears that the 5 measures do not capture the exact same thing, i.e. "trust", or equivalently, it seems that the common perception of there being one homogenous kind of trust is inadequate in our context. In the following section we will perform a closer analysis of the potential determinants of trust to achieve an understanding of causal relationships, and therefore the individual meanings of our measures.

Determinants of Suppliers' Trust – IPR Hold-up, Pay and Fairness

Both our model and our industry survey suggest that the inherent hold-up problem is at the center of trust formation: suppliers "sink" effort into research and design for parts which result in blueprints. After obtaining these blueprints, the superior bargaining position of the OEM enables him to extract additional rents. Therefore our favored interpretation of trust frames it as the belief of the supplier regarding the probability that the OEM will refrain from such undesired behavior.

To test this, we turned to the suppliers evaluation of such behavior by the OEM in our questionnaire, specifically the frequency of conflicts regarding the treatment of patents and trade-secrets as well as the frequency with which the OEM passes on technological secrets of the supplier to third parties without permission.

Clearly, there are more direct ways to extract rents from a position of power, especially by exerting pressure in price negotiations. Therefore for pre-development we looked at the degree to which the OEM shares the (considerable) risk of higher than expected costs. For development, we can use the evaluation of the adequacy of license fees in the case that the OEM makes use of protected know-how of the supplier. And for series production, both an evaluation of the frequency with which the OEM demands lump-sum price reductions in renegotiations as well as the extent to which he attempts to extract cost information by employing sub-supplier management were available to us.

We performed OLS-regressions with the individual trust-measures as the dependent and the measures introduced above as explanatory variables, while controlling for the size and technological sophistication of the part. **Table 4.5** below presents the coefficients and p-values, neglecting the effects of the product-type dummies.

	Trust 1	**Trust 2**	**Trust 3**	**Trust 4**	**Trust 5**
Pre-Development					
Frequency IPR conflicts	-.637 (.05)	-.323 (.01)	-.632 (.00)	-.149 (.277)	-.416 (.02)
How often does OEM leak supplier's IPR	-.147 (.05)	-.521 (.00)	-.291 (.00)	-.153 (.00)	-.225 (.00)
OEM shares risk of higher development costs	.180 (.06)	-.156 (.31)	.040 (.82)	.242 (.00)	.301 (.00)
Development					
Frequency IPR conflicts	-.302 (.08)	-.463 (.00)	-.450 (.00)	-.118 (.11)	-.089 (.39)
How often does OEM leak supplier's IPR	-.170 (.17)	-.392 (.00)	-.333 (.00)	-.134 (.01)	-.116 (.09)
Adequacy of license fees	-.035 (.74)	.100 (.32)	.398 (.00)	.134 (.04)	.256 (.00)
Series Production					
Frequency price re-negotiation (lump sum)	.007 (.91)	-.124 (.13)	-.079 (.45)	-.180 (.00)	-.265 (.00)
Efforts of OEM to extract cost information	-.026 (.65)	-.042 (.41)	-.038 (.54)	-.151 (.00)	-.166 (.00)

Table 4.5.: Determinants of Trust measures.
Coefficients of OLS regressions controlling for product type and (p-values) reported.

The pattern that emerges lends itself to interpretation: The first trust-measure (*importance of trust for the supplier to initiate pre-development co-operation with OEM*) is not significantly influenced by any of the answers to the questions aiming at compensation and pay. Instead, there is a very strong negative correlation (-.637) with the reported frequency of IPR conflicts during pre-development, i.e. the more frequent IPR conflicts, the lower this measure of trust. A much weaker, but still significant effect with the expected sign (-.147) results from the OEM leaking sensitive IP-related data more often. IPR-conflicts during the development phase have a relatively strong negative effect as well (-.302), the p-value of .08 is in part explained by the smaller number of common observations.

The second trust-measure (*mutual trust with respect to IPR*) follows the same general pattern, though with different individual weightings. Again, the compensation measures show no significant influence on the trust-measure. But here, the leaking of sensitive information shows a far stronger effect (-.521 in pre-development, -.392 in development) than in the former case. Further, in this measure the importance of IPR-conflicts during the development phase (-.463) is higher than during the pre-development phase (-.323), which probably reflects the fact that a higher share of respondents to this question were involved in later development stages.

The picture changes for the remaining trust measures three through five (*Importance of trust for the OEM's supplier choice* for pre-development, development and series production, respectively). For measures four and five, the importance of IPR-related behavior decreases markedly, while the effects of adequate compensation and price-cutting become significant with the expected signs across the board. Interestingly, the third trust measure, related to pre-development presents itself as a hybrid case, in which the IPR-related factors are still predominant, but nevertheless also the adequacy of license fees plays an important role (.398), while the other compensation related measures do not. The results further provide a judgment on the – as of yet unanswered – issue of *whose* trust these questions truly measure: As perceived misbehavior by the OEM affects them in a significantly negative manner, we feel comfortable using them as measures for the supplier's trust in the OEM.

To crudely summarize these findings: Trust 1 and Trust 2 are negatively associated with attempts at rent extraction by the OEM in the area of IPR with slightly different focuses, while they are not affected by direct attempts at price-reductions. This balance shifts toward the latter for Trust 3 and even more so for the measures Trust 4 and Trust 5. After going to great lengths to establish these measures, we now try to show that the underlying construct significantly affects relevant economic behavior and outcomes.

4.3.4. Mistrust and Underinvestment

The main hypothesis derived from our model states that higher levels of trust should be associated with more relationship specific investment by suppliers. Due to the cross-sectional structure of our data-set, determining the direction of causality is an issue. The case for higher investment by suppliers leading to higher levels of trust of suppliers in the OEM can be made: Less investment may lead to more conflicts between the parties, which lets trust deteriorate. While acknowledging this, we find the opposite argument more convincing, but leave the issue open and prefer to refer to association instead of causality.

Measuring investment of suppliers poses a challenge. As we do not observe a direct measure, we propose two proxies related to the quality of parts instead. One standard interpretation of quality related effort in the literature is that it affects failure rates of parts (see, for example, Taylor and Wiggins (1997)). Along these lines, we asked respondents two questions: *With respect to the part considered, how often do quality problems occur?* and ... *how often do recall actions occur during series production?* Both questions are measured on a 5-point scale with 1 resembling the lowest and 5 the highest frequency. The correlation between them is (only) 0.41, which can be explained by the fact that not every quality problem leads to a recall and the latter are extremely infrequent. As a consequence, 89% of respondents reported a 1–very infrequently for the recall question, which poses serious empirical limitations.[28] Specifically, in 50% of the cases in which the lowest possible value for recalls is reported, there is some higher level of quality problems. On the other hand, when quality problems become more frequent/severe (levels of 3 or higher), which in itself is rather rare (17 cases), in 54% of cases respondents report levels of recall frequency above 1.

In general, difficulties arise when trying to assess underinvestment-related quality issues empirically, as a) the observed failure rates of cars cannot necessarily be linked to individual parts, b) the diligence of the manufacturer in assembly also affects quality and c) if quality problems are diagnosed before the parts are installed, this is generally not observable. The huge advantage of our questionnaire is that the responses are part-specific which address issue a). The phrasing of the question address issue c). By including customer- or OEM-effects in the regressions, we hope to alleviate issue b). A potential drawback is the fact that the frequencies are self-reported, so that respondents may be tempted to under-report problems. To counter this, complete anonymity was guaranteed at the outset of the study.

For our empirical strategy, we choose the following approach with y denoting the frequency of problems arising, κ denoting a constant, i denoting the part in question and j denoting the customer:

$$y_{ij} = \kappa + \alpha_j + \beta * trust_{ij} + \gamma * dummy\_soph_i + \delta * dummy\_big_i + \epsilon_{ij} \tag{4.6}$$

This **Model 1** spells out the set of OLS-regressions including customer fixed-effects (α).[29] To address the issue of subjective differences in the understanding of the questions at hand, we also specified dummies which take the value 1 whenever the answer to the question is the lowest possible frequency (i.e. 1) and

[28] The model quality, as represented by the F-statistic, in general is very low as meaningful variation is absent.

[29] We also perform (more suitable to the data at hand) tobit and ordered-logit regressions, which delivered qualitatively identical results. We do not report the results in this paper, but they can be provided by the authors upon request.

take the value 0 whenever the reported frequency is larger than this. This allows us to estimate a probit-model, in which y is simply replaced by the probability that no problems occur/problems occur as rarely as possible (**Model 2**). This specification has the further advantage that the results can be more readily interpreted. We would expect negative coefficients for β in Model 1 and positive coefficients for β in Model 2 from our hypothesis.[30]

We estimate both models with and without customer fixed-effects. We include these to capture, for example, potential complementary effort exerted by the customers (OEMs), which also may influence the probability of quality issues arising. As this effort may affect suppliers' incentives to provide better quality, it is not clear how these effects are directed. As a consequence of the very limited variation in the frequency of recalls, we only present the set of OLS regressions for this question to show that the trust measures contribute explanatory power in the cases that recalls are reported. The results of the three sets of regressions can be found in **Tables B.8, B.9** and **B.10**.

[**Table B.8, B.9** and **B.10** about here]

First note that the coefficients in all estimations have the expected signs, though their levels of significance vary. The latter can, to a large degree, be explained by the fact that different numbers of observations are available for the individual regressions. In those including Trust 4, which include the highest number of observations (122), β is significantly different from zero at the 5% level for all specifications. As one would expect, both technological sophistication and size (or interfaces within) raise the probability of quality issues significantly across the board. The size of the coefficients is noteworthy – more complex parts that are both technologically sophisticated and large are more than 50% more likely to have quality issues. The other surprising result is the extent to which our measure Trust 4 influences the quality level as reported by suppliers. An increase in the measure by 1 (i.e. 1.1 standard deviations)

[30] As Model 1 makes use of more variation within the outcome variable, it uses the existing information more efficiently.

decreases the probability of quality issues arising by 12.9% (excluding costomer fixed effects) or even 16.7% (including customer fixed effects), see **Table B.9**. We take this as evidence in favor of our first central hypothesis.

4.3.5. Trust and Sourcing Decisions

In this context, the underlying causal mechanism is closely related to the hypothesis one would expect: If the level of competition affects trust, in the sense that the nature of arms-length market agreements is less amenable to fostering trust relationships, then one should be more inclined to side with the existing literature in expecting a negative association. On the other hand, we could also find ourselves in a setup similar to the model we propose. Then, the supplier faces attempts at rent extraction by the OEM *either* through exploitation of the hold-up situation *or* through competition. In this case more competition may be a signal of higher levels of compliance with property rights by the party with the superior bargaining power, and therefore higher levels of trust. Therefore, at best we can determine correlation with our data.

Our empirical approach is to analyze how three different measures of supplier competition, i.e. the number of parallel suppliers contracted to work on pre-development, development and series production projects, respectively, is associated with our measures of trust derived above, while controlling for the size and technological complexity of the product in question. As a second test, we also directly check for the relationship between the number of competitors and the methods of rent extraction that we identified in **Section 4.3.2** as affecting the relevant measures of supplier trust most. For pre-development, this is the frequency of IPR conflicts and the frequency with which the OEM has discontinued development projects in the past 5 years.[31] For development, this is the frequency of IPR conflicts and the adequacy of license fees, while for

[31] We observe multi-collinearity between IPR conflicts and the OEM leaking supplier IPR (the same holds for development), therefore we omit the latter from the regression. If we use the risk of higher development costs, we loose all but 50 observations. Therefore we replace this variable with the frequency of discontinuation of projects, which captures the opposite of the effect of risk-sharing.

series production, this is the frequency of price re-negotiation and the OEM's efforts at extracting cost information.

We again use simple OLS, as we are mainly interested in correlation between the variables. In the following approach, y is the number of suppliers employed by the OEM, i is the part in question and j is the customer, κ again denotes a constant. x resembles either a trust measure or the vector of the two rent-extraction mechanism described above, depending on the specification.

$$y_{ij} = \kappa + \beta * x_{ij} + \gamma * dummy\_soph_i + \delta * dummy\_big_i + \epsilon_{ij} \tag{4.7}$$

Regarding our second hypothesis, we would expect a positive sign for β in the specifications with the trust measures. Higher levels of trust would be associated with a higher incentive to extract rents through inducing more competition, i.e. employing more parallel suppliers. For the specifications with the direct measures of rent-extraction, we would expect the opposite: The more an OEM, for example, exerts price pressure directly on the supplier, the less he needs to induce competition through multiple parallel suppliers.

[**Table B.11, B.12, B.13** and **B.14** about here]

The results of the regression on the trust measures for the number of suppliers involved in pre-development are interesting in that even the significance of the model itself is rejected (**Table B.11**). There appears to be no significant correlation between our trust measures and the number of suppliers in pre-development, with the same also holding for the complexity and size of the product. The "direct" approach on the other hand yields interesting insights into the posed question (see column 1 in **Table B.14**): We find a highly *negative* correlation between the frequency of conflicts regarding IPR and the number of suppliers involved in the pre-development project. This points in the direction of our hypothesis, in that the supplier's intellectual property plays a pivotal role in the course of pre-development, and abuse of his rights in this realm is a strong indication of attempts at rent extraction. The fact that there is no significant correlation with the size and the complexity of the part may be attributable to the fact that pre-development is by definition not model specific – e.g. research on the application of a new alloy can be connected to a car-model and part ex post, but at the time the basic research takes place, this is far from clear.

With respect to the development stage, the models including the trust-measures directly related to the development process (Trust 2 and Trust 3) have the strongest explanatory power (**Table B.12**). While the coefficients here are not significant, they are positive and in the case of Trust 2 very close to significance with a p-value of .16, which points in the direction of our hypothesis. For the development phase, we also find overal negative and significant effects of the size and complexity of the product in question on the number of suppliers involved in the process. We may capture the effect of less competition in these more differentiated good markets, here. Finally, the direct methods of rent extraction in the case of development show neither a significant positive nor negative sign (see column 2 in **Table B.14**).

In series production, the evidence is mixed (**Table B.13**). Interestingly, the coefficient for Trust 5, the measure directly related to series production, is insignificant (p=.758). Amongst the others, Trust 3 shows a strong positive correlation with the number of suppliers involved in the first stage of production, while Trust 2 also has a positive (even larger) coefficient but is barely insignificant at the 10% level (p=.132). Only the coefficient of Trust 1 is out of line, which resembles the importance of trust for the supplier approaching the OEM for a pre-development project. The "direct" approach yields a very clear picture (see column 3 in **Table B.14**). The frequency of price renegotiations with lump sum reductions demanded by the OEM is highly significantly and *negatively* associated with the number of suppliers involved in the beginning of series production. In our view this strongly indicates that these two instruments of rent extraction are substitutive.

Overall, we find somewhat mixed evidence, with most of it in support of our second hypothesis. We actually consider our findings to be relatively straightforward, if one allows for the fact that we used interview data, in which each observations is compiled from multiple participants and the questions included psychologically loaded terms such as "trust". The structure of our data does not allow us to completely disentangle these ideas, but we strongly believe that more research into this is warranted and important.

4.4. Conclusion and Outlook

Trust is an important ingredient in almost all meaningful social and economic interactions. While most recent economic research, to a large extent due to availability of data, has focused on the willingness of individuals to trust others in general, we were able to shed light on the role of trust that has developed – i.e. been fostered or squandered – in specific economic relationships. In this context, we propose an intuitive definition for trust that we demonstrate in a formal model: In the presence of a hold-up situation, we define trust as the subjective probability that the party with the superior bargaining situation will be exploitative.

We then show in the empirical part of our study that behavior that can be construed as appropriating rents in excess of the formal property rights of the OEM does lead to lower levels of supplier trust, using various different measures for this. We further show that "trust" is not a single homogenous construct, but instead different people in the same firm will focus on different factors when asserting their level of trust in the other party. Some of our measures are more closely related to the adherence to and respect for intellectual property rights, other measures are associated with fair compensation vs. frequent price renegotiations.

We find that an OEM's investment in his supplier's trust by forgoing these (often short-term) opportunities at appropriating rent can pay off: In relationships signified by higher levels of trust we find indication for significantly higher investment by the supplier resulting in fewer failures and callbacks.

Our model makes a second prediction which may appear more surprising at first glance: We show that inducing upstream competition is a substitute for the OEM exploiting an existing hold-up situation – our model predicts that higher levels of trust of the supplier in the OEM should be associated with a *larger* number of upstream suppliers competing. In contrast to results from the literature on relational contracts our empirical findings support this to a certain degree, though not unequivocally. While employing a larger number of suppliers is related with significantly lower levels of rent extraction through other channels, such as price renegotiations ex post, the effects of our trust measures on the number of suppliers is mixed. In the face of our finding that there is no one-size-fits-all concept of "trust", this was to be expected – and shows that more specific research in this direction could be very useful.

Finally, we believe that our understanding of trust as being connected to hold-up situations or similar settings is extremely useful, but clearly does not cover close to all potential applications. In particular, we believe that an application to settings of incomplete contracts, potentially also in procurement environments such as Bajari and Tadelis (2001), could greatly enhance our understanding in this field.

A. Appendix to Chapter 3 - Regression Results

Dep. Variable: Prob(Severance)	Model 1a	Model 2a	Model 3a
Own music of severer at $t = 0$	**.0000139***** (2.72e-06)	-	**.0000169***** (2.65e-06)
Own music of severed at $t = 0$	**-.0000182*** (7.71e-06)	-	**-.0000161**** (7.10e-06)
Play-lists of severer at $t = 0$	**0.00017***** (0.0000301)	-	.0000592 (.000043)
Play-lists of severed at $t = 0$	.0000641 (.000183)	-	.000207 (.0000197)
Music uploaded by severer per day	-	**.00483***** (.00053)	**.00136***** (.000382)
Music uploaded by severed per day	-	**-.0120*** (.00634)	**-.00966*** (.00555)
Play-lists created by severer per day	-	**-.0449***** (.00652)	**-.0186***** (.00465)
Play-lists created by severed per day	-	.0775 (.055)	.0754 (.0467)
observations	1,521	1,521	1,521
pseudo R^2	.429	.312	.469
Chi^2	347.12	252.54	379.64

Table A.1.: Free Riding and Negative Reciprocity a): Logistic regressions on the probability that a given friendship is severed.
Marginal effects and standard errors at mean reported: * significant at 10%, ** significant at 5%, *** significant at 1%

Dep. Variable: Prob(Severance)	Model 1b	Model 2b
Songs listened per day by severer	**-.130*** (.009)	-.0450 (.0455)
Songs listened per day by severed	**.00952***** (.00244)	**.00519***** (.00154)
Own music of severer at $t = 0$	-	**.0000145***** (2.38e-06)
Own music of severed at $t = 0$	-	**-.0000215*** (.0000115)
Play-lists of severer at $t = 0$	-	.0000623 (.0000451)
Play-lists of severed at $t = 0$	-	.0003717 (.000234)
Music uploaded by severer per day	-	**.000746***** (.000286)
Music uploaded by severed per day	-	**-.00862**** (.00366)
Play-lists created by severer per day	-	**-.0279***** (.00707)
Play-lists created by severed per day	-	**.0901**** (.00707)
observations	1,477	1,477
pseudo R^2	.0473	.358
Chi^2	26.64	201.58

Table A.2.: Free Riding and Negative Reciprocity b): Logistic regressions on the probability that a given friendship is severed.
Marginal effects and standard errors at mean reported: * significant at 10%, ** significant at 5%, *** significant at 1%

Dep. Variable: music files added	Model 1	Model 2	Model 3
music friends	-5.44e-06	-.0000134	-.0000292
	(9.66e-06)	(.0000138)	(.0000199)
own music	-.0602665	-.0666011	-.0920428
	(.0585873)	(.0827902)	(.0928703)
Δmusic friends	**4.24e-06***	**5.56e-06****	**6.75e-06****
	(2.29e-06)	(2.49e-06)	(2.91e-06)
experienced dummy	-	-.0423019	.4954036
		(.0632985)	(.5668963)
time	-	.0234988	.0301578
		(.0155074)	(.016886)
comments	-	1.478684	1.352071
		(4.517221)	(4.346862)
messages	-	.0867128	-.1812929
		(1.362083)	(1.33175)
play-lists friends	-	-	.0073593
			(.0072321)
own playlists	-	-	.4954036
			(.5668963)
constant	3.364962	3.906991	5.863737
	(2.266137)	(2.804955)	(3.610217)
observations	177,828	177,828	177,828
groups	25,031	25,031	25,031
overall R^2	.0294	.0241	.0208
F	2.96	5.52	9.15
rho	.924	.933	.962

Table A.3.: Free Riding and Positive Reciprocity for Information Provision, FE-regressions controlling for potential user-cluster heteroskedasticity, entire sample

* significant at 10%, ** significant at 5%, *** significant at 1%

Dep. Variable: music files added	Model 1	Model 2	Model 3
music friends	-.0000124	-.0000293	**-.0000572*** * (.0000345)
	(.0000198)	(.0000197)	(.0000345)
own music	-.0593032	-.0661795	-.1171941
	(.0590668)	(.0836462)	(.1071663)
Δmusic friends	3.87e-06	**9.61e-06****	**.0000105****
	(2.81e-06)	(3.71e-06e-06)	(4.06e-06)
experienced dummy	-	-1.636376	-1.314448
		(2.713547)	(2.485865)
time	-	.3370503	.2793232
		(.17308)	(.1722218)
comments	-	1.48062	1.352071
		(4.511961)	(4.346862)
messages	-	.0533741	-.1812929
		(1.376997)	(1.33175)
play-lists friends	-	-	.0126578
			(.0123971)
own playlists	-	-	1.011472
			(1.176967)
constant	11.12651	5.51823	14.21924
	(7.765481)	(9.963463)	(13.35873)
observations	24,413	24,413	24,413
groups	9,428	9,428	9,428
overall R^2	.0472	.0409	.0317
F	2.27	5.93	6.67
rho	.725	.757	.896

Table A.4.: Free Riding and Positive Reciprocity for Information Provision, active user sub-sample
* significant at 10%, ** significant at 5%, *** significant at 1%

Dep. Variable: play-lists added	Model 1	Model 2	Model 3
play-lists friends	-.0002687	-.0007544	**-.0013166****
	(.0005835)	(.0007107)	(.0006179)
own play-lists	**-.1671713*****	**-.2086505*****	**-.2901718*****
	(.0393094)	(.0453584)	(.0489084)
Δplay-lists friends	**.0009297****	**.0010383****	**.0010542*****
	(.000391)	(.0004011)	(.0003998)
experienced dummy	-	.0046399	.0072075
		(.0084659)	(.0085653)
time	-	**-.0046172*****	.0019167
		.0010961	(.002046)
comments	-	-.0715414	-.0329556
		(.4122087)	(.3545851)
messages	-	.1556359	.0831986
		(.1184353)	(.1029976)
dummy active	-	-.0000785	.0040834
		(.0079719)	(.0073418)
music friends	-	-	**-5.21e-06*****
			(1.97e-06)
own music	-	-	**.0123697****
			(.0056172)
constant	.3138434	.4337222	.5619726
	(.0710501)	(.0549475)	(.2757704)
observations	177,828	177,828	177,828
groups	25,031	25,031	25,031
overall R^2	.0435	.0000	.0005
F	9.43	16.51	22.27
rho	.797	.923	.975

Table A.5.: Free Riding and Positive Reciprocity for Information Organization, entire sample
* significant at 10%, ** significant at 5%, *** significant at 1%

Dep. Variable: play-lists added	Model 1	Model 2	Model 3
play-lists friends	-.0003837	-.0006297	-.0012703
	(.0007793)	(.0009335)	(.0009274)
own play-lists	**-.123411*****	**-.1525017*****	**-.2550607*****
	(.029512)	(.0836462)	(.0808259)
Δplay-lists friends	**.0012131****	**.0012077****	**.0012261****
	(.0005418)	(.0005203)	(.0005189)
experienced dummy	-	.0961424	-.0053575
		(.1195154)	(.1435919)
time	-	-.0199061	.0036746
		(.0116495)	(.011876)
comments	-	-.0874014	-.0446645
		(.3545696)	(.3362418)
messages	-	.0924774	.0663375
		(.1224277)	(.1116464)
music friends	-	-	**-.0000168*****
			(3.25e-06)
own music	-	-	**.0108558***
			(.0064666)
constant	.5901143	.979095	1.46338
	(.1505223)	(.353222)	(.9223387)
observations	24,413	24,413	24,413
groups	9,428	9,428	9,428
overall R^2	.0892	.0026	.0008
F	12.18	11.53	24.35
rho	.524	.651	.914

Table A.6.: Free Riding and Positive Reciprocity for Information Organization, active user sub-sample
* significant at 10%, ** significant at 5%, *** significant at 1%

B. Appendix to Chapter 4: Descriptive Statistics and Regression Results

Variable	Mean (Std. Dev.)	Min	Max	Obs.
When is supplier asked to participate?	2.77 (1.37)	1	6	144
How often is progress coordinated?	2.98 (.57)	1	5	151
Share of efforts absorbed by supplier	3.50 (1.33)	1	5	142
Cost reimbursement if subsequent contract	2.31 (1.52)	1	5	246
Cost reimbursement if no subsequent contract	2.39 (1.59)	1	5	232
Specificity development objectives wrt...				
... content	2.33 (.97)	1	5	350
... time-frame	1.85 (.96)	1	5	350
... financial engagement	2.22 (1.14)	1	5	343
OEM's supplier choice criteria:				
... importance of supplier price	5.10 (1.16)	1	6	158
... importance of duration cooperation	4.70 (.99)	1	6	160
... importance of trust	4.89 (.98)	1	6	159

Table B.1.: Relationship Characteristics: **Pre-Development** (Suppliers' view)

Variable	Mean (Std. Dev.)	Min	Max	Obs.
How specific and detailed are specifications?	2.39 (1.02)	1	5	231
Supplier's degree of freedom	2.91 (.86)	1	5	231
Desired degree of freedom	3.62 (.77)	1	5	229
OEM's contribution to development	2.37 (1.10)	1	5	200
Frequency of IPR conflicts	2.24 (.87)	1	5	194
OEM's supplier choice criteria:				
... importance of supplier price	5.37 (.72)	2.5	6	387
... importance of duration cooperation	4.52 (1.00)	1	6	387
... importance of personal contact	4.52 (.98)	1	6	387
... importance of certification	4.39 (1.14)	1	6	377
... importance of trust	4.90 (.93)	1	6	384

Table B.2.: Relationship Characteristics: **Development** (Suppliers' view)

Variable	Mean (Std. Dev.)	Min	Max	Obs.
How often does OEM produce part himself?	1.69 (1.31)	1	6	210
OEM's supplier choice criteria:				
... importance of supplier price	5.70 (.52)	3	6	253
... importance of duration cooperation	4.38 (1.07)	1	6	253
... importance of personal contact	4.44 (1.10)	1	6	253
... importance of certification	4.28 (1.19)	1	6	250
... importance of trust	4.73 (.98)	1	6	252

Table B.3.: Relationship Characteristics: **Series Production** (Suppliers' view)

Variable	Mean (Std. Dev.)	Min	Max	Obs.
Number of competing suppliers	2.16 (.84)	1	5	137
Frequency of subsequent development projects	3.23 (1.11)	1	5	322
How often were projects discontinued in last 5 yrs.	2.00 (.88)	1	5	139
How often were the following employed...				
... preselection of a specific supplier	4.43 (1.26)	1	6	351
... procurement among a ltd. number of suppliers	3.95 (1.44)	1	6	338

Table B.4.: Procurement Decisions: **Pre-Development** (Suppliers' view)

Variable	Mean (Std. Dev.)	Min	Max	Obs.
Frequency joint procurement dev. and production	3.76 (1.24)	1	5	363
Number of suppliers employed during...				
... product planning	2.22 (1.13)	1	5	167
... product specification	2.03 (1.02)	1	5	177
... concept development	2.12 (1.07)	1	5	208
... detailed development	1.51 (0.90)	1	5	210
How often were the following employed...				
... preselection of a specific supplier	3.06 (1.52)	1	6	259
... procurement among a ltd. number of suppliers	5.18 (1.10)	1	6	264
... open procurement	1.97 (1.41)	1	6	255

Table B.5.: Procurement Decisions: **Development** (Suppliers' view)

Variable	Mean (Std. Dev.)	Min	Max	Obs.
Number of suppliers employed...				
... at production start	1.20 (.58)	1	5	251
... after 1-2 years	1.47 (.78)	1	5	249
... after more than 2 years	1.59 (.81)	1	5	246
How often were the following employed...				
... preselection of a specific supplier	2.98 (1.63)	1	6	248
... procurement among a ltd. number of suppliers	4.55 (1.52)	1	6	248
... open procurement	2.44 (1.66)	1	6	243

Table B.6.: Procurement Decisions: **Series Production** (Suppliers' view)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) Trust 1	1.00								
(2) Trust 2	**.250**	1.00							
	(.048)								
(3) Provision IPR by supplier	-.0296	-.037	1.00						
	(.764)	(.663)							
(4) Provision IPR by OEM	-.036	**.202**	**.443**	1.00					
	(.718)	(.016)	(.000)						
(5) Difference in the Provision IPR	-.009	**-.185**	**.756**	**-.254**	1.00				
	(.930)	(.029)	(.000)	(.000)					
(6) Provision Costs by supplier	.037	.005	.120	.032	.121	1.00			
	(.557)	(.968)	(.229)	(.748)	(.225)				
(7) Trust 3	**.535**	**.432**	-.046	**.136**	**-.147**	-.018	1.00		
	(.000)	(.000)	(.569)	(.092)	(.071)	(.889)			
(8) Trust 4	**.339**	**.385**	-.030	.052	-.078	**-.193**	**.509**	1.00	
	(.000)	(.000)	(.663)	(.447)	(.260)	(.002)	(.000)		
(9) Trust 5	**.320**	**.382**	-.090	-.011	-.104	**-.143**	**.408**	**.700**	1.00
	(.000)	(.002)	(.365)	(.912)	(.296)	(.025)	(.001)	(.000)	

Table B.7.: Pairwise Correlations of Trust Measures
(p-values)

Trust 1: importance of the trust relationship with the OEM in a firm's decision to initiate a pre-development project (six-point scale ranging from 1–no relevance to 6–very high relevance). Trust 2: How do you evaluate mutual trust between OEM and supplier with respect to honoring each otherÕs intellectual property rights (IPR)? (five-point scale ranging from 1–very little to 5–very high) Trust 3 to 5: How important is trust in the OEM's selection of a supplier in pre-develpoment (Trust 3), development (Trust 4), series production (Trust 5)? (six-point scale ranging from 1–no relevance to 6–very high relevance)

	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**	**(10)**
Trust 1	-.052	-.054								
	(.395)	(.42)								
Trust 2			-.072	-.069						
			(.396)	(.51)						
Trust 3					**-.120***	-.099				
					(.087)	(.192)				
Trust 4							**-.133****	**-.159****		
							(.03)	(.022)		
Trust 5									-.066	-.080
									(.437)	(.440)
d_soph	**.298****	**.327****	.085	.144	.049	.084	**.244****	**.271****	**.354*****	**.401*****
	(.026)	(.020)	(.614)	(.418)	(.766)	(.625)	(.024)	(.015)	(.001)	(.006)
d_big	.189	.223	.204	.135	.251	.277	**.218***	**.215***	.183	.186
	(.171)	(.138)	(.26)	(.364)	(.152)	(.146)	(.058)	(.076)	(.193)	(.23)
OEM-FE	**no**	**yes**	**no**	**yes**	**no**	**yes**	**no**	**yes**	**no**	**yes**
# obs.	95	95	67	67	70	70	122	122	90	90
R^2	.102	.162	.05	.171	.092	.238	.118	.162	.126	.187

Table B.8.: OLS-regression results (Model 1) for frequency of quality problems
coefficients and (p-values) reported: * significant at 10%, ** significant at 5%, *** significant at 1%

	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**	**(10)**
Trust 1	.041	.040								
	(.427)	(.487)								
Trust 2			.085	.101						
			(.212)	(.262)						
Trust 3					.096	.111				
					(.111)	(.128)				
Trust 4							**.129****	**.167****		
							(.026)	(.016)		
Trust 5									.086	.092
									(.249)	(.308)
d_soph	**-.240****	**-.283****	-.028	-.084	-.030	.024	**-.207****	**-.237****	**-.298*****	**-.349*****
	(.028)	(.014)	(.836)	(.574)	(.833)	(.881)	(.029)	(.016)	(.008)	(.003)
d_big	**-.224****	**-.278****	**-.278***	-.242.	**-.344****	**-.478*****	**-.246****	**-.258****.	**-.219***	**-.238***
	(.050)	(.027)	(.051)	(.162)	(.015)	(.007)	(.015)	(.015)	(.064)	(.065)
OEM-FE	**no**	**yes**	**no**	**yes**	**no**	**yes**	**no**	**yes**	**no**	**yes**
# obs.	95	95	67	67	70	70	122	122	90	90
Ps-R^2	.096	.149	.074	.186	.103	.241	.107	.138	.127	.169

Table B.9.: Probit-regression results (Model 2), probability of not observing quality problems average marginal effects and (p-values) reported: * significant at 10%, ** significant at 5%, *** significant at 1%

	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**	**(10)**
Trust 1	**-.071****	**-.090****								
	(.033)	(.011)								
Trust 2			-.052	-.087						
			(.290)	(.168)						
Trust 3					**-.081****	**-.090****				
					(.027)	(.028)				
Trust 4							-.049	**-.073****		
							(.142)	(.049)		
Trust 5									-.063	**-.116***
									(.204)	(.055)
d_soph	.041	.047	.067	.078	-.018	-021.	.052	.062	.080	.070
	(.568)	(.523)	(.508)	(.471)	(.84)	(.821)	(.389)	(.317)	(.319)	(.398)
d_big	.109	.089	.014	-.054	.080	.060	.087	.061	.075	.040
	(.148)	(.260)	(.898)	(.659)	(.392)	(.560)	(.168)	(.347)	(.360)	(.65)
OEM-FE	**no**	**yes**	**no**	**yes**	**no**	**yes**	**no**	**yes**	**no**	**yes**
# obs.	92	92	64	64	66	66	115	115	87	87
R^2	.077	.195	.033	.155	.089	.210	.048	.118	.053	.153

Table B.10.: OLS-regression results (Model 1) for frequency of product-related recalls coefficients and (p-values) reported: * significant at 10%, ** significant at 5%, *** significant at 1%

	(1)	(2)	(3)	(4)	(5)
Trust 1	-.006				
	(.947)				
Trust 2		-.132			
		(.123)			
Trust 3			-.098		
			(.214)		
Trust 4				-.084	
				(.322)	
Trust 5					.208
					(.157)
d_soph	-.006	.059	.099	.062	-.035
	(.982)	(.695)	(.522)	(.686)	(.891)
d_big	-.023	.133	.075	.119	.001
	(.934)	(.446)	(.68)	(.504)	(.997)
# obs.	62	126	134	127	61
R^2	.000	.028	.020	.014	.036

Table B.11.: OLS-regression results for number of suppliers **pre-development** coefficients and (p-values) reported: * significant at 10%, ** significant at 5%, *** significant at 1%

	(1)	(2)	(3)	(4)	(5)
Trust 1	-.018				
	(.820)				
Trust 2		.125			
		(.161)			
Trust 3			.079		
			(.333)		
Trust 4				.101	
				(.188)	
Trust 5					-.139
					(.276)
d_soph	**-.382***	-.239	**-.299***	**-.238***	**-.380***
	(.079)	(.126)	(.068)	(.060)	(.079)
d_big	-.161	**-.458****	**-.488****	-.208	-.162
	(.498)	(.015)	(.014)	(.126)	(.494)
# obs.	97	113	119	206	96
R^2	.052	.116	.117	.036	.066

Table B.12.: OLS-regression results for number of suppliers **development** coefficients and (p-values) reported: * significant at 10%, ** significant at 5%, *** significant at 1%

	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Trust 1	**-.061****				
	(.049)				
Trust 2		.200			
		(.132)			
Trust 3			**.188****		
			(.047)		
Trust 4				.041	
				(.260)	
Trust 5					-.031
					(.758)
d_soph	**-.187****	**-.448***	**-.462***	**-.196*****	**-.186****
	(.014)	(.077)	(.057)	(.010)	(.014)
d_big	**-.153***	-.256	-.229	**-.154****	**-.167****
	(.053)	(.323	(.364)	(.049)	(.034)
# obs.	244	59	63	249	248
R^2	.066	.170	.195	.057	.053

Table B.13.: OLS-regression results for number of suppliers **series production** coefficients and (p-values) reported: * significant at 10%, ** significant at 5%, *** significant at 1%

	(1)	(2)	(3)
Frequency IPR conflicts pre-development	**-.389*** **		
	(.003)		
Frequency discont. of projects by OEM	.079		
	(.350)		
Frequency IPR conflicts development		.035	
		(.710)	
Adequate license fees		-.014	
		(.879)	
Frequency price renegotiation			**-.125*** **
			(.000)
Efforts at extraction of cost info			.005
			(.801)
d_soph	.028	**-.265***	-.090
	(.139)	(.091)	(.192)
d_big	.261	-.079	-.090
	(.179)	(.643)	(.200)
# obs.	111	147	195
R^2	.102	.024	.133

Table B.14.: OLS-regression results: Relationship of different rent-extraction devices and number of suppliers during **(1) pre development**, **(2) development** and **(3) series production**
coefficients and (p-values) reported: * significant at 10%, ** significant at 5%, *** significant at 1%

Bibliography

Aghion, P., Y. Algan, P. Cahuc, and A. Shleifer (2008). Regulation and distrust. *Harvard University, mimeo.*.

Aghion, P., M. Dewatripont, and P. Rey (2002). On partial contracting. *European Economic Review 46*, 745–753.

Akerlof, G. A. and J. L. Yellen (1990). The fair wage-effort hypothesis and unemployment. *The Quartely Journal of Economics 105*, 255–283.

Andreoni, J. (1988). Why free ride? *Journal of Public Economics 37*, 291–304.

Anton, J. J. and D. A. Yao (2007). Finding llostpprofits: An equilibrium analysis of patent infringement damages. *The Journal of Law, Economics & Organization 23, 1*, 186–207.

Asvanund, A., K. Clay, R. Krishnan, and M. D. Smith (2004, June). An empirical analysis of network externalities in peer-to-peer music-sharing networks. *Information Systems Research 15*(2), 155–174.

Bajari, P. and S. Tadelis (2001). Incentives versus transaction costs: a theory of procurement contracts. *RAND Journal of Economics 32*, 387–407.

Baker, G., R. Gibbons, and K. J. Murphy (1994). Subjective performance measures in optimal incentive contracts. *The Quarterly Journal od Economics 109*(4), 1125–1156.

Baker, G., R. Gibbons, and K. J. Murphy (1999). Informal authority in organizations. *The Journal of Law, Economics, & Organization 15*, 56–73.

Bala, V. and S. Goyal (2000, September). A noncooperative model of network formation. *Econometrica 68*(5), 1181–1229.

Ben-Shahar, O. and J. J. White (2006). Boilerplate and economic power in auto manufacturing contracts. *Michigan Law Review 104*, 953–82.

Bessen, J. E. and M. J. Meurer (2006). Patent litigation with endogenous disputes. *The American Economic Review 96*, 77–81.

Bessen, J. E. and M. J. Meurer (2008a). *Patent Failure.* Princeton University Press.

Bessen, J. E. and M. J. Meurer (2008b, February). The private costs of patent litigation. *mimeo*.

Birke, D. (2008). The economics of networks - a survey of the empirical literature. *Nottingham University Occasional Paper Series*, 1–33.

Bloch, F. and B. Dutta (2009). Communication networks with endogenous link strength. *Games and Economic Behavior 66*, 39–56.

Boldrin, M. and D. K. Levine (2002). The case against intellectual property. *American Economic Review 92*, 209–212.

Boldrin, M. and D. K. Levine (2008). *Against Intellectual Monopoly.* Cambridge University Press.

Bottazzi, L., M. Da Rin, and T. Hellmann (2007). The importance of trust for investment: Evidence from venture capital. *mimeo*.

Bottazzi, L., M. Da Rin, and T. Hellmann (2009). What is the role of legal systems in financial intermediation? theory and evidence. *Journal of Financial Intermediation 18*, 559–598.

Bramoullé, Y. and R. Kranton (2007, July). Public goods in networks. *Journal of Economic Theory 135*, 478–494.

Brown, M., A. Falk, and E. Fehr (2004). Relational contracts and the nature of market interactions. *Econometrica 72*(3), 747–780.

Bull, C. (1987). The existence of self-enforcing implicit contracts. *The Quartely Journal of Economics 102*, 147–159.

Butler, J., P. Giuliano, and L. Guiso (2009). The right amount of trust. *NBER Working Paper No. 15344*.

Cagno, D. D. and E. Sciubba (2007). *Advances in Decision Making Under Risk and Uncertainty*, Chapter The Determinants of Individual Behaviour in Network Formation: Some Experimental Evidence, pp. 219–241. Springer-Verlag Berlin Heidelberg.

Cagno, D. D. and E. Sciubba (2008, January). Social networks and trust: not the experimental evidence you may expect. working paper.

Caillaud, B. and A. Duchene (2006). Patent office in innovation policy: Nobody's perfect. *mimeo*.

Calzolari, G. and G. Spagnolo (2009). Relational contracts and competitive screening. *CEPR Discussion Paper No. 7434*.

Chiou, J.-Y. (2008). The patent quality control process: Can we afford an (rationally) ignorant patent office? *mimeo*.

Cho, M. (2006, December). Endogenous formation of networks for local public goods. working paper.

Choi, J. P. (2006). How reasonable is the 'reasonable' royalty rate? damage rules and probabilistic intellectual property rights. *CESifo Working Paper No. 1778*.

Coleman, J., E. Katz, and H. Menzel (1957). The diffusion of an innovation among physicians. *Sociometry 20*, 253–270.

Cook, J. P. (2007). On understanding the increase in u.s. patent litigation. *American Law and Economics Review 9*(1), 48–71.

Corbae, D. and J. Duffy (2008). Experiments with network formation. *Games and Economic Behavior 64*, 81–120.

Crampes, C. and C. Langinier (2002). Litigation and settlement in patent infringement cases. *The RAND Journal of Economics 33*(2), 258–74.

Falk, A. and M. Kosfeld (2003, May). It's all about connections: Evidence on network formation. *IZA Discussion Paper* (777).

Farrell, J. and C. Shapiro (2008). How strong are weak patents? *American Economic Review 98*, 1347–1369.

Fehr, E. (2009). On the economics and biology of trust. *Journal of the European Economic Association 7*, 235–266.

Fehr, E. and S. Gächter (2000). Fairness and retaliation: The economics of reciprocity. *Journal of Economic Perspectives 14*, 159–181.

Goeree, J. K., A. Riedl, and A. Ule (2008). In search of stars: Network formation among heterogeneous agents. *Games and Economic Behavior article in press.*

Gowrisankaran, G. and J. Stavins (2004). Network externalities and technology adopotion: Lessons from electronic payments. *The RAND Journal of Economics 35*, 260–276.

Goyal, S. (2005). Strong and weak links. *Journal of the European Economic Association 3*, 608–606.

Goyal, S., M. J. v. d. Leij, and J. L. Moraga-González (2006). Economics: An emerging small world. *Journal of Political Economy 114*, 403–412.

Granovetter, M. S. (1974). *Getting a Job: A Study of Contacts and Careers.* Harvard University Press.

Granovetter, M. S. (1976). Network sampling: Some first steps. *The American Journal of Sociology 81*, 1287–1303.

Grossman, S. J. and O. D. Hart (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *The Journal of Political Economy 94*, 691–719.

Gu, B., Y. Huang, W. Duan, and A. B. Whinston (2007, October). Online peer-to-peer communities: An empirical investigation of a music sharing network as a dynamic two-sided network. working paper.

Gu, B., Y. Huang, W. Duan, and A. B. Whinston (2009). Indirect reciprocity in online social networks - a longitudinal analysis of individual contributions and peer enforcement in a peer-to-peer music sharing network. *McCombs Research Paper Series No. IROM-06-09*, 1–40.

Guiso, L., P. Sapienza, and L. Zingales (2004). The role of social capital in financial development. *American Economic Review 94*, 526–556.

Guiso, L., P. Sapienza, and L. Zingales (2005). Trusting the stock market. *NBER Working Paper 11648*.

Guiso, L., P. Sapienza, and L. Zingales (2006). Does culture affect economic outcomes? *Journal of Economic Perspectives 20*, 23–48.

Guiso, L., P. Sapienza, and L. Zingales (2009). Cultural biases in economic exchange? *The Quarterly Journal of Economics 124*, 1095–1131.

Hall, B. H. and D. Harhoff (2004). Post-grant reviews in the u.s. patent system - design choices and expected impact. *Berkeley Technology Law Journal 19*(3), 989–1015.

Harman, H. H. (1976). *Modern factor analysis* (3rd ed.). The University of Chicago Press.

Hart, O. (2008). Hold-up, asset ownership, and reference points. *mimeo*.

Hart, O. and J. Moore (1988). Incomplete contracts and renegotiation. *Econometrica 56*, 755–785.

Hart, O. and J. Moore (1990). Property rights and the nature of the firm. *Journal of Political Economy 98*, 1119–1158.

Hart, O. and J. Moore (2007). Incomplete contracts and ownership: Some new thoughts. *American Economic Review 97*, 182–186.

Häussler, C., D. Harhoff, and E. Müller (2009). To be financed or not... - the role of patents for venture capital financing. *CEPR Discussion Paper* (7115).

Jackson, M. O. (2004). *Group Formation in Economics; Networks, Clubs and Coalitions*, Chapter A Survey of Models of Network Formation: Stability and Efficiency. Cambridge University Press.

Jackson, M. O. and A. Watts (2002). The evolution of social and economic networks. *Journal of Economic Theory 106*, 265–295.

Jackson, M. O. and A. Wolinsky (1996). A strategic model of social and economic networks. *Journal of Economic Theory 71*, 44–74.

Kosfeld, M. (2003). Network experiments. *IERE Working Paper* (152).

Kossinets, G. and D. J. Watts (2006). Empirical analysis of an evolving social network. *Science 311*, 88–90.

Kumar, R., J. Novak, and A. Tomkins (2006). Structure and evolution of online social networks. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining, Philadelphia 2006*, pp. 611–617.

La Porta, R., F. L. de Silanes, A. Shleifer, and R. W. Vishny (1997). Trust in large organizations. *American Economic Review 87*(2), 333–338.

Lemley, M. A. (2001). Rational ignorance at the patent office. *Northwestern University Law Review 95*, 1495–1529.

Lemley, M. A. and C. Shapiro (2005). Probabilistic patents. *The Journal of Economic Perspectives 19*, 75–98.

Levin, J. (2003). Relational incentive contracts. *American Economic Review 93*(3), 835–857.

MacCallum, R. C., K. F. Widaman, S. Zhang, and S. Hong (1999). Sample size n factor analysis. *Psychological Methods 4*, 84–99.

MacLeod, W. B. (2007). Reputations, relationships and contract enforcement. *Journal of Economic Literature 45*, 595–628.

Meurer, M. J. (1989, Spring). The settlement of patent litigation. *The RAND Joural of Economics 20*(1), 77–91.

Müller, F., K. O. Stahl, and F. Wachtler (2008). Upstream relationships in the automotive industry: A contractual perspective. *mimeo*.

Newman, M. (2001). The structure of scientific collaboration networks. *Proceedings of the national Academy of Sciences of the United States of America 98*, 404–409.

O'Dea, D. (2008, November). Heterogeneity in social networks and the classical theory of public goods. working paper.

Prady, D. (2008). ïs that a gun in your pocket... ?ädvertising and evaluation. *mimeo*.

Reitzig, M., J. Henkel, and C. Heath (2007). On sharks, trolls, and their patent prey - unrealistic damage awards and firms' strategies of "being infringed". *Research Policy 36*, 134–54.

Rosenblat, T. S. and M. M. Mobius (2004). Getting closer or drifting apart? *Quarterly Journal of Economics 119*, 971–1009.

Saloner, G. and A. Shepard (1995). Adoption of technologies with network effects: An empirical examination of the adoption of automated teller machines. *The RAND Journal of Economics 26*, 479–501.

Sapienza, P., A. Toldra, and L. Zingales (2007). Understanding trust. *NBER WP 13387*.

Schankerman, M. and S. Scotchmer (2001). Damages and injunctions in protecting intellectual property. *The RAND Journal of Economics 32*(1), 119–220.

Schuett, F. (2009). Inventors and impostors: An economic analysis of patent examination. *mimeo*.

Shapiro, C. (2003). Antitrust limits to patent settlements. *The RAND Journal of Economics 34*, 391–411.

Straus, J. and N.-S. Klunker (2007). Harmonisation of international patent law. *International Review of Intellectual Property and Competition Law 8*, 907–36.

Taylor, C. R. and S. N. Wiggins (1997). Competition or compensation: Supplier incentives under the american and japanese subcontracting systems. *The American Economic Review 87*, 598–618.

Womack, J. P., D. T. Jones, and D. Roos (1991). *The Machine That Changed the World: The Story of Lean Production.* HarperPaperbacks.

Xia, M., Y. Huang, W. Duan, and A. B. Whinston (2007). To share or not to share? - an empirical analysis on user decisions in online sharing communities. working paper.

Eidesstattliche Erklärung

Hiermit erkläre ich, die vorliegende Dissertation selbständig angefertigt und mich keiner anderen als der in ihr angegebenen Hilfsmittel bedient zu haben. Insbesondere sind sämtliche Zitate aus anderen Quellen als solche gekennzeichnet und mit Quellenangaben versehen.

Mannheim, 12.05.2010.

Johannes Koenen

Curriculum Vitae

1981	Born in Frankfurt am Main, Germany
1998	US High-School Diploma, Phillips Academy Andover, Andover, Massachusetts
2000	Graduation from Secondary School, Kaiserin-Friedrich-Gymnasium, Bad Homburg, Germany
2000-2006	Undergraduate Studies in Economics, University of Mannheim, Germany
2003-2006	Undergraduate Studies in Law, University of Mannheim, Germany
2006-2010	Graduate Studies in Economics at the Center for Doctoral Studies in Economics (CDSE), University of Mannheim, Germany
2007-2010	Research Assistant at the Chair for Applied Microeconomics, Professor Konrad Stahl, PhD, University of Mannheim, Germany
2009	Visiting Scholar at Dartmouth College, Hanover, New Hampshire