

Mang, Constantin

**Research Report**

## Market Consequences of ICT Innovations

ifo Beiträge zur Wirtschaftsforschung, No. 70

**Provided in Cooperation with:**

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

*Suggested Citation:* Mang, Constantin (2016) : Market Consequences of ICT Innovations, ifo Beiträge zur Wirtschaftsforschung, No. 70, ISBN 978-3-95942-017-4, ifo Institut - Leibniz-Institut für Wirtschaftsforschung an der Universität München, München

This Version is available at:

<https://hdl.handle.net/10419/167367>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.


Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*



70

## ifo Beiträge zur Wirtschaftsforschung

### **Market Consequences of ICT Innovations**

Constantin Mang

**ifo** Institut

Leibniz-Institut für Wirtschaftsforschung  
an der Universität München e.V.

Herausgeber der Reihe: Clemens Fuest  
Schriftleitung: Chang Woon Nam

70

ifo Beiträge  
zur Wirtschaftsforschung

## **Market Consequences of ICT Innovations**

Constantin Mang

## **Bibliografische Information der Deutschen Nationalbibliothek**

Die Deutsche Nationalbibliothek verzeichnet diese Publikation  
in der Deutschen Nationalbibliografie; detaillierte bibliografische  
Daten sind im Internet über  
<http://dnb.d-nb.de>  
abrufbar

ISBN-13: 978-3-95942-017-4

Alle Rechte, insbesondere das der Übersetzung in fremde Sprachen, vorbehalten.  
Ohne ausdrückliche Genehmigung des Verlags ist es auch nicht gestattet, dieses  
Buch oder Teile daraus auf photomechanischem Wege (Photokopie, Mikrokopie)  
oder auf andere Art zu vervielfältigen.

© ifo Institut, München 2016

Druck: ifo Institut, München

ifo Institut im Internet:  
<http://www.cesifo-group.de>

# Preface

This volume was prepared by Constantin Mang while he was working at the Ifo Institute. It was completed in 2014 and accepted as a doctoral thesis by the Department of Economics at the University of Munich. It includes a short introduction and four self-contained chapters that focus on ICT innovations in different markets.

*Chapter 1* gives a short introduction into the topic and lays out some general aspects about the methodological frameworks used in this volume. *Chapter 2* is focused on the effects of broadband Internet on the housing market in order to quantify the value of broadband access in terms of a premium on housing rents. Methodologically, the chapter uses micro data from Germany's largest online platform for real estate advertisements and estimates a hedonic model with rent prices. *Chapter 3* investigates the association of online job search and matching quality using individual-level data from the German Socio-Economic Panel (SOEP). It measures matching quality by the respondent's evaluation of his new job compared to his former job. The results show that job changers who found their new job online are better matched than their counterparts who found their new job through traditional channels. *Chapter 4* provides evidence on the effectiveness of PC use in schools across countries. Using data from the international student achievement test TIMSS, it estimates the effect of PC use on student test scores and finds that using a PC for some activities has positive effects on student achievement, whereas using a PC for other activities has negative effects. *Chapter 5* investigates the effects of a mobile Internet infrastructure upgrade on the take-up of location-based online services. Using data from the largest German online platform for restaurant reviews, it finds that upgrading the mobile Internet infrastructure increases the number of restaurant reviews and the share of reviews written on smartphones.

Keywords: Broadband Internet, mobile Internet, value, hedonic pricing, job search, matching quality, PC use in classrooms, e-learning, infrastructure upgrade.

JEL-No.: L86, L96, R31, R42, H54, O18, I21, I28, J64.



## **Acknowledgements**

First and foremost, I would like to thank my supervisor Ludger Woessmann for his constant guidance, advice and support during my PhD studies. I am also deeply grateful to Oliver Falck, my co-supervisor. His helpful comments, inspiring ideas and continuous encouragements were an enormous help to me. Furthermore, I thank Monika Schnitzer for joining the dissertation committee.

I sincerely thank Oliver Falck and Ludger Woessmann for their collaboration on Chapter 4 and Martin Micheli for his collaboration on Chapter 2 of this thesis. I also owe many thanks to my current and former colleagues at the Ifo Center for the Economics of Education and Innovation for inspiring conversations and helpful suggestions. Particularly, I would like to thank Nadine Fabritz for supporting me in any possible way during the finalization phase of this thesis. I have greatly benefitted from working together with Francesco Cinnirella, Nina Czernich, Bernhard Enzi, Alexandra Heimisch, Erik Hornung, Stefan Kipar, Franziska Kugler, Philipp Lergetporer, Susanne Link, Andreas Mazat, Natalie Obergruber, Marc Piopiunik, Jens Ruhose, Benedikt Siegler, Martin Schlotter, Ruth Maria Schüler, Guido Schwerdt, Benedikt Siegler and Katharina Werner. I also received comments and suggestions from many other friends and colleagues that unfortunately cannot all be mention here. I gratefully acknowledge financial support from Deutsche Telekom AG.

I owe my deepest gratitude to my parents and my sister Isabelle for their unconditional support that was the basis of my work on the PhD thesis.





# **Market Consequences of ICT Innovations**

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

2014

vorgelegt von  
**CONSTANTIN MANG**

Referent: Prof. Dr. Ludger Wößmann  
Korreferent: Prof. Dr. Oliver Falck  
Promotionsabschlussberatung: 5. November 2014



---

# Contents

---

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Consequences for Markets . . . . .	2
1.2	Methodological Framework . . . . .	6
1.3	Outline of the Thesis . . . . .	8
<b>2</b>	<b>The Value of Broadband Internet - Evidence from the Housing Market</b>	<b>13</b>
2.1	Related Literature . . . . .	16
2.2	A Hedonic Market for Broadband Internet . . . . .	19
2.3	Data . . . . .	23
2.4	Empirical Strategy . . . . .	26
2.4.1	Estimation of the Hedonic Price for Broadband . . . . .	26
2.4.2	Historical Features of Broadband Infrastructure in Germany . . . . .	27
2.4.3	Exploiting Variation in the Distance Between Properties and MDFs . . . . .	29
2.4.4	Exploiting Variation in the Roll-out Costs for Properties that are Far Away from their MDF . . . . .	31
2.4.5	Exploiting Variation from a Technological "Mistake" in Eastern Germany . . . . .	32
2.5	Results . . . . .	33
2.5.1	The Effect of DSL Availability on Property Prices . . . . .	33
2.5.2	Interpretation of the Estimates Size . . . . .	36
2.5.3	Comparison of our Estimates with Estimates from Previous Studies . . . . .	38
2.6	Validation of Results with Survey Data from SOEP . . . . .	39
2.6.1	Description of SOEP Data . . . . .	39
2.6.2	The Value of Broadband for SOEP Households . . . . .	40
2.7	Conclusion . . . . .	41
	Figures and Tables . . . . .	43
	Appendix . . . . .	52
<b>3</b>	<b>Online Job Search and Matching Quality</b>	<b>57</b>
3.1	Data and Methodology . . . . .	60
3.1.1	Data on Job Search Methods and Matching Quality . . . . .	60
3.1.2	Estimation Strategy . . . . .	62
3.2	The Association of Internet Job Search and Matching Quality . . . . .	64
3.2.1	Baseline Results . . . . .	64
3.2.2	Results in Matched Samples . . . . .	65
3.2.3	Effect Heterogeneity . . . . .	66
3.2.4	How Internet Job Search Compares to Other Search Methods . . . . .	69
3.3	Discussion of Potential Selection Issues . . . . .	70
3.3.1	Systematic Differences in Job Seeker Characteristics . . . . .	70
3.3.2	Selection into Search Intensities . . . . .	72

3.3.3	Selection of Advertised Jobs . . . . .	74
3.4	Conclusion . . . . .	76
	Figures and Tables . . . . .	77
<b>4</b>	<b>PC use in Classrooms and Student Achievement</b>	<b>91</b>
4.1	Related Literature . . . . .	95
4.2	Data . . . . .	97
4.3	Empirical Framework . . . . .	102
4.4	Results . . . . .	107
4.5	Student Heterogeneities . . . . .	110
4.6	Country Heterogeneities . . . . .	113
4.7	The Fourth-Grade Sample . . . . .	119
4.8	Conclusion . . . . .	122
	Figures and Tables . . . . .	125
	Appendix . . . . .	136
<b>5</b>	<b>The Effect of Mobile Internet Infrastructure on Local Online Services</b>	<b>141</b>
5.1	Related Literature . . . . .	144
5.2	Data . . . . .	146
5.2.1	Restaurant and Review Data from Qype . . . . .	147
5.2.2	Spatial and Regional Data . . . . .	148
5.2.3	Descriptive Statistics . . . . .	150
5.3	Identification Strategy . . . . .	152
5.4	Effects on the Number of Restaurant Reviews . . . . .	157
5.5	Mobile Reviews and the Restaurant Rating . . . . .	160
5.5.1	Do Mobile Reviewers Give Better Ratings? . . . . .	161
5.5.2	Are Restaurants with More Reviews Better Rated? . . . . .	164
5.6	Robustness . . . . .	165
5.6.1	Threats to the Exogeneity of a Clear Line-of-Sight . . . . .	165
5.6.2	Sensitivity to Different Sample Restrictions . . . . .	168
5.7	Conclusions . . . . .	170
	Figures and Tables . . . . .	172
	<b>References</b>	<b>190</b>

---

## List of Tables

---

2.1	Descriptive statistics of ImmoScout24 data in the full sample and three estimation samples . . . . .	46
2.2	OLS and IV estimates in the distance sample of ImmoScout24 properties . . . . .	47
2.3	OLS and IV estimates in the above-threshold sample of ImmoScout24 properties . . . . .	48
2.4	OLS and IV estimates in the OPAL sample of ImmoScout24 properties . . . . .	49
2.5	Comparison of OLS estimates for each sample in east and west Germany . . . . .	50
2.6	OLS and IV estimates for west and east Germany in the SOEP sample . . . . .	51
A2.1	Descriptive statistics of SOEP 2008 survey data on rented properties . . . . .	54
A2.2	The association between Internet use intensity and distance to the MDF . . . . .	55
3.1	All job changers and job changers who found their job through the Internet . . . . .	81
3.2	Sample means by job search method . . . . .	82
3.3	The association between online job search and several matching outcome variables . . . . .	83
3.4	Regression-adjusted matching samples . . . . .	84
3.5	Effect heterogeneity with the dependent variable "skill use" . . . . .	85
3.6	Different job search channels compared to the employment office . . . . .	86
3.7	Reduced sample for comparison of Internet with newspaper and friends . . . . .	87
3.8	Internet job finders and personality traits . . . . .	88
3.9	Controlling for respondent's assessment of working time as a way to account for bias from systematic differences in job changer characteristics . . . . .	89
3.10	Controlling for whether job changers were actively searching for a job . . . . .	90
4.1	PC availability and usage intensities by subject for eighth-grade students in TIMSS . . . . .	127
4.2	Cross-sectional OLS and seemingly unrelated regressions for eighth-grade students . . . . .	128
4.3	Within-student between-subject effect of PC use activities for eighth-grade students . . . . .	129
4.4	The effect of PC use activities in subsamples of top-performing and of disadvantaged eighth-grade students . . . . .	130
4.5	The effect of PC use activities for eighth-grade students in subsamples of countries with specific characteristics . . . . .	131
4.6	PC availability and usage intensities by subject for the full and the same-teacher sample of fourth-grade students . . . . .	132
4.7	Within-student between-subject effect of PC use activities for fourth-grade students . . . . .	133
4.8	The effect of PC use activities in subsamples of top-performing and of disadvantaged fourth-grade students . . . . .	134
4.9	The effect of PC use activities for fourth-grade students in subsamples of countries with specific characteristics . . . . .	135
A4.1	Share of eighth-grade students with between-subject variation in PC use intensities . . . . .	138
A4.2	Share of fourth-grade students with between-subject variation in PC use intensities . . . . .	139
5.1	Descriptive statistics in the full and in the rural sample . . . . .	178

5.2	Effects on the total number of reviews . . . . .	179
5.3	Effects on having at least one mobile review . . . . .	180
5.4	Effects on the share of mobile reviews out of all reviews . . . . .	181
5.5	The source of reviews and their corresponding star rating . . . . .	182
5.6	The source and star rating of reviews with restaurant and author fixed effects . . . . .	183
5.7	Mobile phone brands and the star rating of reviews . . . . .	184
5.8	The number of mobile and non-mobile reviews and the overall star rating of a restaurant	185
5.9	Effects when only the oldest 50 percent of antennas are used to determine visibility .	186
5.10	Effects in samples with different rural area definitions . . . . .	187
5.11	Effects for a restricted sample of restaurants that had at least one review in 2007 . . .	188
5.12	Effects on the total number of reviews when 2007 is used as the base year in difference- in-difference estimations . . . . .	189

---

# List of Figures

---

1.1	The importants of different job search methods over time . . . . .	3
1.2	Students enrolled in HarvardX MOOCs by location as of June 22, 2014 . . . . .	5
2.1	The equilibrium hedonic pricing schedule with bid and offer functions . . . . .	43
2.2	Representation of the "local loop" between households and MDF . . . . .	43
2.3	Screenshot of a property advertisement on ImmobilienScout24.de . . . . .	44
2.4	The average broadband penetration of municipalities in 30 quantile bins . . . . .	45
A2.1	The availability of fixed-line broadband access with at least 16 Mbps over the availability of broadband with at least 1 Mbps in relation to the average distance between municipality centroids and MDFs on the state level in 2010 . . . . .	52
A2.2	Share of SOEP households with DSL in different distance bins . . . . .	53
A2.3	Fixed-line broadband Internet subscribers in Germany (in million) . . . . .	53
3.1	Distribution of propensity scores for individuals who found their job online and individuals who found their job through other channels . . . . .	77
3.2	Preference of job searchers for electronic and postal applications in Germany over time . . . . .	78
3.3	Usage of CV databases and online career networks for passive job search in Germany	78
3.4	Company types that advertise on two major German online job boards . . . . .	79
3.5	Percentage of search results that match the search request on Monster.de . . . . .	79
3.6	Quality of applications in response to online and print job advertisements for "Head of Procurement" . . . . .	80
4.1	Effect sizes of combined PC use (to process data, practice skills & look up ideas) correlated with GDP per capita and total spending on education . . . . .	125
4.2	Effect sizes of combined PC use (to process data, practice skills & look up ideas) correlated with computer and broadband access . . . . .	126
A4.1	Correlations of GDP per capita and total spending on education with the percentage of 8th grade students using e-learning at least once a month . . . . .	136
A4.2	Correlations of computer and broadband access by country with the percentage of 8th grade students using e-learning at least once a month . . . . .	137
5.1	Screenshots from the Qype iPhone App . . . . .	172
5.2	All restaurants in Germany that can be found on Qype . . . . .	173
5.3	Old (green) and new (red) antenna locations in Germany . . . . .	174
5.4	Digital Surface Model (DSM) and antenna locations . . . . .	175
5.5	Disaggregated population density in the area around Munich . . . . .	175
5.6	Schematic representation of 2G and 3G penetration properties . . . . .	176
5.7	Data volumes in the German mobile network in million GB . . . . .	176
5.8	Number of reviews by visibility . . . . .	177





# Chapter 1

---

## Introduction

---

Information and Communication Technologies (ICT) affect the way markets function. By lowering search and transaction costs, they facilitate market access, increase market transparency and improve the matching between market participants. Today, there is hardly any market that is unaffected by ICT innovations. For example, E-Commerce websites like Amazon.com use the capabilities of ICT to hold a large inventory of products at a centralized location and send them to consumers around the world. Sellers on traditional retail markets are unable to offer an equally large range of products. The market for used goods was traditionally limited to second-hand shops, flea markets and transactions between family and friends. Since the advent of the Internet, eBay and other online platforms are able to match millions of consumers with each other. Search engines and price comparison websites allow consumers to quickly obtain information on products and prices that would either be very costly or impossible to collect without the help of ICT.

The primary role of information technologies in markets is to facilitate communication, coordination and information processing. Due to their wide use, their ongoing technical improvement, and their applications in many different markets, they can be best described as "general purpose technologies" (Bresnahan and Trajtenberg, 1995). These kind of technologies do not only have a direct effect on application sectors and markets, but also enable complimentary innovations in other markets. For example, the Internet has a direct effect on the market for DVD rentals by providing consumers with easier access to movie trailers and ratings. As a general purpose technology, the Internet also enables complimentary innovations and new business models, such as on-demand online video streaming.

One feature that distinguishes many ICT innovations from previous general purpose technologies such as the steam engine or electricity, is that they can spread very quickly. Especially Internet related innovations are able to rapidly disseminate around the world. The main reason for this is that on the Internet, there is no time to manufacture, no inventory management, and no shipping delay (Varian, 2010). Since many Internet innovations consist of bits and bytes, their single components can be easily copied and recombined to create new innovations in different markets. This recombination of ICT enabled innovations leads to what is often described as "recombinant growth" (Weitzman, 1998).


## 1.1 Consequences for Markets

In many ways, the Internet can be considered as a platform that facilitates the exchange of information and reduces transaction costs. Since many market participants can use the same platform at very low costs, the Internet is a valuable tool for matching processes. If we consider matching processes in a network of market participants, an important role of the Internet is to facilitate the formation of the network. This means that more market participants are connected with each other and the probability of one participant ending up without any connection decreases. The role of the Internet for the clearing of the network is less obvious. Under some circumstances, the increased number of connections can make matching more difficult if there is no appropriate clearing mechanism in place.

One example that illustrates the difficulties of network clearing is the housing market. Online platforms for real estate advertisements, such as the one described in Chapter 2, make it easier and less costly for renters to find properties and to contact landlords. Especially in areas with high demand for housing, this leads to many potential renters for every property. Landlords are then faced with the decision of whom they want to let the property to. However, they do not know which other properties potential renters are interested in. Therefore, they might choose a renter that already has offers for other properties, while other potential renters do not get any offer.

Another market where matching plays an important role is the labor market. When job seekers apply for jobs, a network of applicants and firms arises (Gautier and Holzner, 2011). Online job boards, social networks and other Internet resources facilitate the formation of such a network. The wider selection of job advertisements, better search options and lower costs of application lead to more links between job seekers and firms. As a result, more job seekers will be able to establish

Figure 1.1  
The importance of different job search methods over time



Based on own calculations with data from the German Socio-economic Panel (SOEP).

at least one connection to a firm and fewer firms will not receive any application. However, it might become more difficult for firms to select one applicant and for job seekers to accept one offer. It is therefore unclear, whether the Internet makes job matching necessarily more efficient. Less ambiguously, the increased number of connection within the network can lead to better matches between job seekers and firms. Chapter 3 provides evidence that job seekers who found their job through the Internet are better matched than job seekers who used traditional channels like newspapers or friends.

Figure 1.1 shows how the importance of the Internet for matching on the job market has increased over time. The graph depicts how people who recently started a new job found out about this job. The most important source for job seekers are family and friends. Until about 10 years ago, the second most important job search channel was newspapers. However, the importance of newspapers has drastically declined over time and since 2010, the Internet has become the second most important job search channel. Kroft and Pope (2014) provide evidence that the declining role of classified job advertisements in newspapers is a direct consequence of the increasing importance of online job search.

Besides its effects on matching, the Internet also improves the transparency of markets. By lowering the costs of acquiring information, consumers can make more informed purchasing decisions. A classic example are websites that compare prices from different sellers. By connecting

consumers with a large number of sellers, they can be viewed as facilitators of a network of market participants. Especially for homogenous goods, these websites can lead to consumers becoming very price-sensitive (Ellison and Ellison, 2009). The reduction of search costs can increase competition and lead to significantly lower prices (Goolsbee and Brown, 2002). Consumers can also use these online tools for arbitrage across markets. For example, consumers can find out that an airline offers a flight at different prices, depending on the booking location.


Another aspect of the increased market transparency that is induced by the Internet is information on the reputation of a seller. Traditionally, consumers had a relatively small network of sellers they were connected to. The small number of connections led to many repeated transactions which helped to build trust between buyers and sellers. In the Internet, the potential number of connections is drastically higher and it might therefore be more difficult to establish trust. The way that many electronic markets alleviate this problem is by providing information on the reputation of a seller. In a way, the idea is to not only connect buyers with sellers, but also to create connections among buyers. One application that has been frequently studied in the empirical literature is the reputation system of eBay (e.g., Jin and Kato, 2006; Resnick et al., 2006; Lucking-Reiley et al., 2007; Cabral and Hortacsu, 2010).

In the offline world, reputation systems are especially important for experience goods. For example, many consumers use online review platforms to see what previous customers have written about a given hotel or restaurant. There is evidence that these reviews have a sizable impact on bookings for hotels (Ye et al., 2011) and restaurants (Anderson and Magruder, 2012). Online user reviews have also been shown to affect other markets, such as the market for movies (Chintagunta, Gopinath and Venkataraman, 2010) or books (Chevalier and Mayzlin, 2006). Chapter 5 uses data from Germany's largest platform for restaurant reviews and shows how these reviews are affected by the availability of mobile broadband Internet.

Another important market development that is caused by ICT innovations is what some people refer to as the "death of distance" (Cairncross, 1997). The argument is that the decrease in transaction costs will ultimately eliminate distance as a cost factor. In some industries, we have seen developments that go into this direction. For example, the distance to bricks-and-mortar retailers has become less important, as the Internet allows consumers to order almost any good to any location. Telecommuting has in some industries replaced the need for workers to physically commute to their office.

One area where the death of distance hypothesis has recently gained traction is education. Universities around the world have started to offer Massive Open Online Courses (MOOCs). These virtual courses are usually free and open to anyone who wants to participate. The design of the courses often has a focus on interactivity. Since the number of students that enrol in MOOCs is too large for conventional student-lecturer interaction, innovative online tools serve as interactive elements. Also the network that is established among MOOC students plays an important role in the course design (Waldrop, 2013).

Figure 1.2  
Students enrolled in HarvardX MOOCs by location as of June 22, 2014



Based on own calculation with data from Nesterko et al. (2014).  
The location of students is determined by their IP address.

In 2012, the Massachusetts Institute of Technology and Harvard University founded a MOOC platform called edX. Today, many other institutions, such as Stanford University and the University of California, Berkeley, have joined the platform. Currently, HarvardX, the MOOCs by Harvard University, has more than 1.33 million enrolled students.<sup>1</sup> Figure 1.2 shows the wide geographic dispersion of HarvardX students. While the majority of students is located in the U.S., 237,505 students are located in Europe and 269,253 students are located in Asia. It is remarkable that the platforms allows thousands of students from developing countries to attend university courses taught by Harvard professors.

<sup>1</sup> Students that are enrolled in more than one course are counted multiple times.

Similar developments take place in primary and secondary education. For example, the Florida Virtual School offers more than 120 courses taught by certified teachers.<sup>2</sup> In Germany, students in some schools on East Frisian islands are taught virtually by teachers on the mainland.<sup>3</sup> However, there is mixed empirical evidence on how the use of computers in classrooms affects student learning. Chapter 4 shows that the effectiveness of PC use in schools depends a lot on the applications that the PCs are used for.

## 1.2 Methodological Framework

In order to provide evidence-based policy advice regarding ICT innovations, it is crucial to identify causal effects rather than mere correlations between a treatment and an outcome variable. For example, one might be tempted to interpret the correlation between broadband Internet and economic growth as causal in the sense that broadband diffusion increases prosperity. However, it could also be the case that growth leads to higher demand for broadband Internet and thereby increases its availability. This is a typical example of reverse causality. Another issue is that there might be a third variable that is correlated with broadband Internet diffusion and also has a positive effect on economic growth. For example, the availability of broadband Internet could be correlated with complementary computer skills in the population. If computer skills also affect economic growth, we would get an omitted variable bias by not taking this effect into account. In order to avoid the endogeneity problems of reverse causality and omitted variable bias, economists use a set of methods that allow for identification of causal effects.<sup>4</sup> These methods have fundamentally changed the way that empirical economics developed over the last decades (Angrist and Pischke, 2010).

A first step towards avoiding bias from omitted variables is to control for confounding factors. The problem with this approach is that we hardly ever observe all potential factors that might influence the outcome variable. If one of these unobservable factors is correlated with the treatment variable, the estimation results remain biased. The same applies to matching methods that aim at finding observations in the treatment and control group that are similar to each other. As long as the unobserved variables remain unobserved, neither conventional regressions nor matching methods can avoid omitted variable bias. However, matching can be used as an adjunct to regressions in order

---

<sup>2</sup> <http://www.flvs.net/students/pages/find-course.aspx>

<sup>3</sup> One example is the Inselschule Borkum, see <http://www.inselschule-borkum.de>.

<sup>4</sup> See Angrist and Pischke (2008) for a detailed discussion of these methods and Schlotter, Schwerdt and Woessmann (2011) for a non-technical overview.

to ensure overlap between treatment and control group. In Chapter 3, I therefore use propensity score matching algorithms to restrict the control groups to those individuals that are most similar to individuals in the treatment group.

One way to tackle omitted variable bias is to use several observations of the same individual in a fixed-effects model. For example, panel data contains repeated observations of individuals over time. By including individual fixed effects, it is possible to avoid estimation bias from time-invariant unobserved characteristics of the individual. The same approach can be used if the same individuals are observed in different settings. In Chapter 4, we observe student test scores in math and science. By comparing the same student in different subjects, we can alleviate omitted variable bias and selection concerns. One important assumption of conventional fixed-effects models is that the effect of the treatment is constant over observations. In order to relax this assumption we estimate a correlated random effects model in Chapter 4.

Another method that is often used to account for time-invariant confounding factors is the difference-in-differences approach. The idea behind this identification strategy is to have a treatment and a control group that are observed before and after a treatment, such as a policy intervention. For example, in Chapter 5 we look at two groups of restaurants that differ in their specific geographic surrounding. Our treatment of interest is the introduction of 3G technology. While restaurants in the first group are likely to benefit from this technology early on, restaurants in the second group are not. Since we observe both groups before and after the treatment, we can estimate the difference over time and between the two groups. This difference-in-differences accounts for unobserved time-invariant heterogeneities. The crucial identifying assumption is that the trends of the two groups would have been the same in absence of the treatment.

All of the above methods have limitations compared to the "gold standard" of causal inference: randomized experiments. By randomly assigning individuals into a treatment and control group, we can be sure that observed differences between the groups can be interpreted as the causal effect of the treatment on the outcome of interest. However, it is usually difficult and often impossible to conduct a controlled experiment. In many cases, a more feasible approach, is to exploit exogenous variation in the treatment variable. In the ideal case, nature or institutional rules separate treatment and control groups as good as randomly. Chapter 2 provides several examples for such quasi-experiments. In order to estimate the effect of broadband Internet availability on rent prices, we exploit that after the German reunification, some parts of eastern Germany received a new telephone technology called OPAL. At that time, OPAL was considered state-of-the-art and the former monopolist for telephone

services wanted to gain experiences with the new technology. Several years later, a broadband Internet technology became popular that was compatible with traditional telephone connections but incompatible with OPAL. For this reason, households connected with OPAL were not able to receive broadband Internet. In this example, the critical assumption is that the selection of OPAL areas by the monopolist was random. If this was the case, our setting would be close to a controlled field experiment, even though no one intended to conduct an experiment on the effect of broadband Internet.

### **1.3 Outline of the Thesis**

This thesis consists of four self-contained empirical studies that investigate the effects of ICT innovations on different markets. Chapters 2 and 3 provide insights on how the Internet affects prices on the housing market and the matching quality of workers and firms on the job market. Chapter 4 analyzes the effect of PC use in schools on student achievement and thereby addresses the education and, after all, the labor market. Chapter 5 investigates the effect of an upgrade in mobile Internet infrastructure on the usage of a local online service that is focused on improving market transparency by providing customer reviews for restaurants. The following paragraphs provide a brief summary of each chapter.

CHAPTER 2 of this thesis is concerned with the value of broadband Internet for consumers. Although the Internet increases consumer surplus in several ways, it is difficult to measure this surplus. A main reason for this difficulty is that many services that can be used online are free of charge. For products without a price, we cannot apply conventional methods to trace out the demand curve and calculate the surplus. Therefore, we apply a different approach and use a unique dataset of real estate advertisements to estimate the hedonic price of broadband Internet. Since broadband availability improves over time due to continuing investments by telecommunication providers and subsidies from municipalities and the state, house sales data are not ideal to estimate the capitalization effect of broadband availability. In order to estimate the value of current broadband availability without making assumptions on its potential improvement over time, we use monthly rent prices instead of sales prices of properties.

Amenities, like broadband, are not randomly distributed across properties. For example, high income areas with a high demand for broadband are likely to get preferential treatment by telecommunication providers. As these areas probably also benefit from other amenities,



such as better schools or cultural offerings, it is difficult to separate the effect of broadband availability from the effect of other amenities. We therefore exploit exogenous variation in broadband availability that goes back to the roots of the voice telephony infrastructure in Germany. When the telecommunication network in Germany was designed, the Internet was not yet existent. The only signals that were sent through copper wires to every household, were analog signals of voice. At the end of the 20th century, when the demand for high speed Internet increased, the DSL technology was developed. The key idea of DSL is to use higher frequency bands to transmit data digitally through the copper wires. The disadvantage of the technology is that the signal loss for high frequencies is higher than for low frequencies. Therefore, some households that are too far away from the next nodal point in the telephone network could not get DSL early on.

In eastern Germany, we exploit another source of exogenous variation in broadband availability. After the German reunification, many parts of the telephone network in eastern Germany had to be replaced. The monopolist for telephone services in Germany at that time decided to roll-out a new technology called OPAL in some eastern German areas. To the best of our knowledge, the selection of these areas was as good as random. Although the OPAL technology was state-of-the-art when it was rolled-out, it turned out to be incompatible with DSL. For this reason, areas that were served by OPAL could not get broadband Internet early on.

Using three different instrumental variables that are based on the exogenous variation in broadband availability, we estimate that properties in municipalities with a 10 percent higher DSL availability have up to 4.7 percent higher rent prices. This corresponds to a monetary value of about €23 per month. However, we find that the effect is very heterogenous across regions. In eastern Germany, we find considerably smaller effects than in western Germany. We also find smaller effects for properties in municipalities that are in the top of the DSL distribution. One interpretation of our results is that broadband Internet is most valuable at the *extensive* margin that determines if broadband is available at all, as opposed to the *intensive* margin that determines the speed of broadband.

CHAPTER 3 investigates the relationship between online job search and the quality of job matches. The Internet has fundamentally changed the way workers and firms are matched on the job market. Compared to newspapers and other traditional job search channels, online job boards presumably lead to better matches by providing a wider choice of job advertisements and more sophisticated methods for finding suitable vacancies. I use survey data from the German Socio-Economic Panel (SOEP) that contains detailed information on how job changers evaluate

their new job compared to their previous job. For example, job changers are asked if they are able to use their skills better in the new job, if they believe to have higher chances of promotion and if they believe to have higher job security. By using information on a new job compared to a previous job of the same person, I am able to estimate a model in the spirit of first-differences.

The results show that job changers who found their new job online are better matched than their counterparts who found their new job through traditional channels. There is evidence that the Internet is an especially valuable job search tool for workers who are distant from the labor market. Job seekers with employment interruptions have significantly better matching outcomes if they find a new job through the Internet. The same holds for job seekers in rural areas. The results hold if I apply propensity score matching methods to restrict the comparison group of offline job seekers to those that are most similar to online job seekers. I also address several selection issues with robustness tests and provide some descriptive evidence to alleviate remaining selection concerns.

CHAPTER 4 provides evidence on the effectiveness of different PC use activities in schools. Although the use of computers in classrooms is often seen as an opportunity for quality improvements, previous studies on the effect of computer use on student achievement are inconclusive. We speculate that one reason for the mixed results that are found in the literature is that the effectiveness of PC use depends to a large extent on the specific activities engaged in on the PC. We use data on eighth-grade students from the international student achievement test TIMSS to show that three distinct PC activities have very different effects on student achievement. Specifically, we analyze how the use of PCs for practicing skills and procedures, processing and analyzing data and looking up ideas and information affects student test scores.

In order to identify the effect of PC use on test scores, we exploit between-subject variation in math and science that allows us to estimate within-student effects, holding unobserved school and student characteristics constant. By exploiting within-student between-subject variation, we also avoid any bias due to student selection into high PC usage schools or into high PC usage classrooms. In conventional first difference models, we implicitly assume that the effect is the same across time, or in our case across subjects. Since this assumption is likely to be violated in our application, we adopt a correlated random effects model that allows us to relax the assumptions of the first-difference estimator.

We find that using a PC to look up ideas and information has positive effects on student achievement in science, whereas using the PC to practice skills and procedures has negative

effects in math and science. The combined effect of all PC use activities is zero in science and negative in math. Although PC use is slightly more beneficial for top-performing students and more detrimental to the achievement of low-performing and non-native students, effect heterogeneities among different groups of students are not very large. We find no evidence that in countries with higher GDP per capita, higher relative educational spending, or better ICT infrastructure PCs are used more effectively by eighth-grade students. We also use data on fourth-grade students that confirm the general pattern we find for eighth-grade students. Since many fourth-grade students are taught by the same teacher in math and science, we are also able to estimate within-student within-teacher effects and thereby avoid bias from unobserved teacher characteristics or selection of teachers into classrooms. Although smaller in size, the estimates in the fourth-grade sample are generally in line with the estimates in the eighth-grade sample.

CHAPTER 5 investigates how an upgrade in mobile Internet infrastructure affects the take-up of location-based online services. Providing for a better broadband Internet infrastructure became an item of high priority on the political agenda of politicians around the world. But network providers are reluctant to upgrade their network as long as the demand does not justify the higher investments. Especially for areas with low initial demand, it is important to know if an infrastructure upgrade affects the usage of the infrastructure. Chapter 5 analyzes the effect of an upgrade in the German mobile Internet infrastructure from a second-generation network (2G) to a third-generation network (3G).

While mobile Internet is useful for a wide variety of applications, such as online search or social networks, it is indispensable for local online services that offer information based on the user's location. Many of these location based services aim at increasing market transparency for consumers by providing reviews on local businesses from previous customers. I use unique data from Germany's largest platform for restaurant reviews to investigate if restaurants in areas where 3G was available early on received more reviews.

To identify the causal effect of the network upgrade, I exploit the geographic location of restaurants relative to cellular network antennas and the terrain around them. Since the 3G signal is much more dependent on a clear line-of-sight than the 2G signal, some restaurants that were previously well covered by 2G are blocked from 3G reception due to geographic obstacles such as hills and forests. Since the actual availability of 3G at a given restaurant cannot be observed, I estimate the reduced-form effect of a clear line-of-sight in a difference-in-difference setting. I find that in areas where for geographical reasons 3G was available early on, restaurants received

more reviews and a larger share of these reviews were written on a smartphone. While I find positive effects of 3G availability on the number and share of mobile reviews, I do not find evidence that reviews submitted from mobile phones are considerably more or less favorable than reviews submitted from PCs.

## Chapter 2

---

# The Value of Broadband Internet - Evidence from the Housing Market\*

---

It is relatively easy to determine the monetary costs of broadband Internet for consumers by looking at broadband subscription prices. For example, an average entry-level broadband subscription costs about \$27 in the U.S., \$16 in South Korea and \$22 in Germany (OECD, 2013).<sup>1</sup> However, these costs underestimate the real value of broadband Internet to consumers. Some of the most frequently used online services, such as E-Mail, search, and social networks, have a cost of zero, even though they are of tremendous value to its users. Due to competition among Internet Service Providers, the prices of Internet access do not necessarily reflect the value of all services that can be used online. One way to derive this value is to ask consumers how much money they would need as a compensation for giving up Internet access. This compensating variation reflects the consumer surplus and therefore the value of Internet access to consumers.<sup>2</sup>

A different way to think about compensating variation is to ask consumers how large a discount in house prices would need to be in order for them to choose a property *without* broadband Internet over an otherwise identical property *with* broadband Internet access. Anecdotal evidence suggests that this discount must be large. In 2012, a survey among 2,000 homebuyers in the U.K. found that broadband Internet is more important to homebuyers than off-street parking, access to shops or a nearby pub. One out of 10 homebuyers indicated to have rejected a potential new home because it

---

\* This chapter was coauthored by Martin Micheli, Rheinisch-Westfälisches Institut für Wirtschaftsforschung.

<sup>1</sup> Prices from September 2012 in USD PPP for an entry-level fixed-line broadband basket with 2 GB data volume and a minimum speed of 0.250 Mbit/s.

<sup>2</sup> Hicks (1942) provides a thorough discussion of relationship between compensating variation and consumer surplus.

had poor Internet connection.<sup>3</sup> A more recent study among users of a large real estate portal showed that information about broadband access at a property's location is more important to users than information about public transport and nearby schools.<sup>4</sup>

This paper uses micro data from Germany's largest online platform for real estate advertisements to estimate a hedonic model with rent prices. We find that a 10 percent increase in broadband availability leads to an increase in rent prices of up to 4.7 percent. This corresponds to a capitalization effect of about €24 per month. Our results indicate that the effects are smaller at the top of the broadband distribution which suggests that the value of broadband is higher at the *extensive* margin than at the *intensive* margin. In eastern Germany, our estimates are considerably lower than in western Germany. We validate our results with survey data from the German Socio-Economic Panel (SOEP), in which we directly observe broadband Internet adoption by households.

In order to identify the causal effect of broadband Internet on rents, we exploit several technical particularities of the decades-old voice telephony network that affect broadband availability in Germany until today. These particularities go back to a time when broadband Internet was non-existent. When broadband was rolled-out in Germany, it was almost exclusively realized through DSL. The DSL technology uses the voice telephony infrastructure, which was designed for the analog transmission of signals, to transmit data digitally. The advantage of DSL is that since it builds on the older infrastructure, roll-out costs are relatively low. The disadvantage is that since the digital signal uses higher frequencies than human voice, the signal loss of DSL is higher than the signal loss of voice telephony. Therefore, some households that can use their telephone lines for regular voice telephony, cannot get DSL. We are using this exogenous variation in DSL availability for several different instrumental variable approaches that lead to comparable results. Since the design of the voice telephony network was determined before the arrival of broadband Internet, the exogenous variation that we exploit cannot be a result of the demand for broadband Internet.

There are many ways in which broadband Internet contributes to consumer surplus.<sup>5</sup> First, the Internet offers numerous possibilities to communicate. Whether E-Mail, instant messaging or social networks, almost all online communication channels can be used for free. Second, access to information is faster, more effective and cheaper on the Internet. Third, many products and services that are offered online do not exist in the offline world. Other products exist offline, but are a lot

---

<sup>3</sup> <http://www.telegraph.co.uk/property/propertynews/9570756/Fast-broadband-more-important-to-house-buyers-than-parking.html>

<sup>4</sup> <http://www.rightmove.co.uk/news/articles/rightmove-news/rightmove-adds-broadband-tool-to-property-details-to-enhance-user-experience>

<sup>5</sup> For a detailed discussion see Morton (2006)

cheaper online. Fourth, the Internet increases market transparency. Many people use websites that compare hundreds of prices for airline tickets or hotels. Additionally, review platforms increase market transparency by allowing consumers to share information about the price and quality of businesses, such as restaurants (Mang, 2014). Fifth, the Internet improves matching processes. Platforms like eBay allow users to buy and sell used goods that would probably not be traded at all in the offline world. The Internet also improves the quality of matches, for example on the labor market (Mang, 2012).

One important way in which the Internet leads to consumer gains is through decreases in prices. The Internet allows many businesses to save costs and operate more efficiently. In competitive markets, this cost reduction will lead to a decrease in prices. Additionally, increased market transparency will foster competition and thereby also lead to lower prices. Empirically, it is well documented that the Internet has led to more competitive markets for cars (Morton, Zettelmeyer and Silva-Risso, 2001), insurances (Goolsbee and Brown, 2002), books (Brynjolfsson, Hu and Smith, 2003), low-cost computer memory modules (Ellison and Ellison, 2009) and many other goods. Besides the increased competition in existing markets, the Internet gives consumers access to a larger variety of products and also leads to the development of completely new goods. Even though their value is hard to assess, new goods are a centerpiece of economic welfare (Bresnahan and Gordon, 1996). While monetary savings due to lower prices are important to consumers, the Internet also leads to significant savings in time. Faster access to information and more efficient ways of communication mean that consumers have more time left for other activities which increases consumer welfare.

Although the Internet increases consumer surplus in many ways, this surplus is difficult to measure. When Robert Solow (1987) famously stated that the computer age can be seen “everywhere except in the productivity statistics,” he put the problem of measuring the value of the information economy into a nutshell.<sup>6</sup> Despite the thousands of new online services that mushroomed in the last two decades, the relative contribution of the information sector (which includes software, broadcasting, telecom, and information and data processing services) to GDP is not larger than it was 25 years ago (Brynjolfsson and Saunders, 2009). This paradox recently revived a public debate about the measurement of the Internet’s value for consumers and the economy. Recent coverage in

---

<sup>6</sup> Although Solow is referring to productivity growth, the problem of measuring “what everyone feels to have been a technological revolution” (Solow, 1987, p. 36) is similar to the problem of measuring the consumer surplus of the Internet.

The Economist,<sup>7</sup> The New York Times<sup>8</sup>, and The Wall Street Journal<sup>9</sup>, illustrate the public interest in this topic.

The remainder of this paper is organized as follows. Section 2.1 provides an overview of the literature on the value of the Internet. Section 2.2 introduces a simple hedonic model for broadband Internet and discusses its implications. Section 3.1 describes the data used in this paper. Section 2.4 presents our identification strategy and gives a short overview of the broadband Internet infrastructure in Germany. Section 3.2 presents our results and Section 3.4 concludes.

## 2.1 Related Literature

Probably the most straightforward way to estimate consumer surplus is to trace out the demand curve. When the price and quality of goods as well as the corresponding demand are known, the only real challenge is to overcome the problem of simultaneity. For the Internet, it is much harder to determine the demand curve and to estimate consumer surplus, because many online services can be used for free. For products without a price, such as E-mail, Facebook or YouTube, the method of estimating consumer surplus by determining the demand curve from observable data cannot be applied. Dutz, Orszag and Willig (2009) and Greenstein and McDevitt (2011*b*) therefore estimate demand and consumer surplus of broadband Internet access, a product that has a price. Greenstein and McDevitt finds that consumer surplus of broadband Internet amounts to \$8.3-\$6.7 billion from 1999 to 2006. One problem of this approach is that the authors are not able to appropriately control for the quality of broadband access over time. Increasing Internet speeds and an increasing number of free services might influence broadband demand independently from its price. While the speed of Internet can be partly accounted for (Greenstein and McDevitt, 2011*a*), many other quality characteristics remain unobserved.

Due to the difficulties of estimating a demand curve for Internet services, many studies try to determine the willingness-of-pay by survey methods. For example, Goolsbee (2006) uses a survey which directly asks consumers how much they are willing to pay for high-speed Internet. This contingent valuation method is often criticized for not measuring real preferences (Diamond and Hausman, 1994). Therefore, other studies implement more elaborate conjoint methods, where

---

<sup>7</sup> <http://www.economist.com/news/finance-and-economics/21573091-how-quantify-gains-internet-has-brought-consumers-net-benefit>

<sup>8</sup> <http://www.nytimes.com/2013/05/01/business/statistics-miss-the-benefits-of-technology.html>

<sup>9</sup> <http://blogs.wsj.com/cio/2013/04/05/beyond-gdp-measuring-value-in-a-service-oriented-information-based-digital-economy>



respondents are presented with different choice sets among which they choose their preferred alternative. For example, Savage and Waldman (2004) give respondents the choice between 64 different Internet access options with different combinations of quality attributes. In a report to the U.S. Federal Communications Commission (FCC), Rosston, Savage and Waldman (2010) present respondents with 8 choice scenarios with different pairs of Internet service alternatives. Although conjoint analyses might be more convincing than conventional contingent methods, the key problem of stated preferences data remains: hypothetical willingness-of-pay is not the same as true economic preferences that are observed on the market.

An interesting alternative to survey-based measures of stated preferences is the use of revealed preference data from experiments. This method has the additional advantage that the functional form assumptions of demand functions can be relaxed (Varian, 1982). Altmann, Rupp and Varaiya (2001) and Varian (2012) use data from an Internet bandwidth experiment in the 1990's. Subjects affiliated with the University of California at Berkeley could choose and switch between various Internet plans with different quality attributes, such as speed and data volume, for different prices. The studies find that subjects place a surprisingly low value to bandwidth. However, the bandwidths used in the experiment lie between 8 and 128 kilobytes per second which is far lower than what we call high-speed Internet today.<sup>10</sup> To our knowledge, there has not been any similar experiment with modern broadband speeds.

A couple of studies focus specifically on free online services. Bughin (2011) combines contingent and conjoint methods to estimate the willingness-to-pay of various online services, such as email, search, social networks and instant messaging. The study also deducts an estimate of the "price" that consumers pay by viewing advertisements during the use of these services. According to Bughin, the highest consumer surplus is created by E-mail, followed by search, social networks and instant messaging. Bughin and Manyika (2014) repeat the analysis with newer data and finds that mobile Internet has significantly increased consumer surplus from free online services between 2010 and 2013.

Varian (2009) takes the viewpoint of advertisers in ad auctions of search engines and calculates the surplus ratio, which he defines as the ratio of aggregate value to aggregate costs. Using data from Google, he finds that this ratio amounts to 2.0-2.3, on average. These values are used by Google to estimate its yearly "economic impact" (e.g., Google Inc., 2012). Chen, Jeon and Kim (2014) look at the consumer side and conduct an experiment to find out how much time subjects save when using

---

<sup>10</sup> In this paper, we define broadband Internet as bandwidths above 1024 kilobytes per second.

online search instead of libraries for information retrieval. Using a random sample of search engine requests, they ask subjects to find answers to the search terms online and offline. On average, online search takes 7 minutes and offline search takes 22 minutes to answer one question. The authors control for result quality and try to mitigate bias from search query selection. Varian (2011) uses these values to estimate consumer surplus based on the average hourly earnings.

The idea of using time for measuring the value of Internet services is also the basis of Goolsbee and Klenow (2006). The authors argue that for goods with low or zero marginal costs and a low share of total costs on spending, consumer surplus can be best estimated by putting the time spent on consuming the good in relation to the opportunity costs of time. The underlying logic is that consumers with higher incomes spend more time using the Internet, an observation that is confirmed by other studies (e.g., Goldfarb and Prince, 2008). Goolsbee and Klenow find that consumer gains from Internet exceed \$3,000 per year for the median consumer. Brynjolfsson and Oh (2012) develop the approach further and allow for interactions of various online and offline activities. The authors also take quality improvements in Internet access into account and only use time use data for free services, such as Facebook, YouTube, Twitter and Wikipedia.

One of the oldest and most popular methods to determine the value of unpriced goods is to use housing market data to estimate hedonic price models. The idea of using house prices to value specific amenities is widely spread at least since the seminal paper of Ridker and Henning (1967). Rosen (1974) provided the first thorough economic interpretation of these models. Although all kinds of unpriced goods have been valued with hedonic models, there is only one study to our knowledge that applies this method to Internet access. Ahlfeldt, Koutroumpis and Valletti (2014) use transaction data from house sales in the U.K. to estimate the value of higher Internet speeds. They find that a 10 percent increase in speed leads to a 1.2 percent increase in house prices. Given the mean house price in their sample, a speed increase of 10 percent corresponds to a monetary value of \$4,296 for the average property. The problem with house sales data is that Internet access technology is constantly upgraded and houses that have limited access to high-speed Internet today will eventually get access to higher Internet speeds in a few years. From the house owner's perspective, the detrimental effect of slow Internet will therefore not necessarily persist over the whole depreciation period of the property. In the end, the house price is therefore also dependent on the expectation of when the property will be connected to high-speed Internet. Taking this expectation into account complicates the calculation of the consumer surplus of high-speed Internet.

## 2.2 A Hedonic Market for Broadband Internet

Rosen (1974) defines the hedonic price as the implicit price of an attribute that is revealed to economic agents from the observed prices of differentiated products and their associated attributes. In our application, we are interested in the implicit price for broadband Internet that is revealed through observed rent prices for properties with and without broadband availability. Following Rosen (1974), we consider a property with characteristics  $Z = (z_1, z_2, \dots, z_n)$ . The price  $P$  of a property is a function of its characteristics  $Z$  which include, for example, the number of rooms, the age, and the neighborhood of the property. We can therefore write:

$$P = P(z_1, z_2, \dots, z_n) \quad (2.1)$$

A buyer (tenant) on the housing market derives this function by comparing the prices of properties with different characteristics. If two properties with identical characteristics are offered for different prices, the buyer will only consider the cheaper property which is why  $P$  gives the minimum price for any combination of property characteristics. We make the reasonable assumption that every consumer rents only one property. The utility of a consumer is then given by  $U(X, z_1, z_2, \dots, z_n)$ , where  $X$  includes all other goods consumed. We set the price of  $X$  equal to unity and measure income  $I$  as units of  $X$ . We assume, all income is spent on  $P$  and  $X$  which gives us the budget constraint  $I - P - X = 0$ . Subject to the budget constraint, consumers will choose  $Z$  and  $X$  such that for the  $i$ th element of  $Z$  the following is satisfied:

$$\frac{\partial P}{\partial z_i} = \frac{\partial U / \partial z_i}{\partial U / \partial X} \quad (2.2)$$

Thus the marginal rate of substitution between  $z_i$  and  $X$  must be equal to the implicit price for  $z_i$  on the housing market. We can describe the optimal bid that a consumer will make for a property with characteristics  $Z$  as a bid function  $\theta$ . Subtracting the bid of a consumer from his income will get the amount that he can spend on all other goods  $X$ . The bid function  $\theta$  must therefore satisfy

$$u = U(I - \theta, z_1, z_2, \dots, z_n) \quad (2.3)$$

Given a fixed utility  $u$  and income  $I$ , a consumer is willing to pay  $\theta(Z; u, I)$  for alternative combinations of  $Z$ .  $\theta_{z_i}$  represents an indifference curve that sets any value of the  $i$ th characteristic of  $Z$  into relation with the amount of  $X$  that is foregone by consuming  $z_i$ . In other words, it reflects the marginal rate of substitution between  $z_i$  and money. Therefore,  $\theta_{z_i}$  can be interpreted as the implicit value that a consumer puts on  $z_i$ , given fixed utility and income.

On the other side of the housing market, sellers (landlords) choose characteristics  $Z$  they want to offer. For example, they can choose to renovate a property, change its layout or install a central heating system. Sellers are also able to influence the Internet speed available at their properties. In many cases, the technical requirements for fast broadband access will be given due to investments by network providers. In other cases, broadband is not available *ex ante* at a property's location. Independently from the *ex ante* availability which is determined by prior investments into the infrastructure, sellers can choose to invest themselves in order to increase Internet speeds. For example, they can spend money on a contract with a network provider that ensures an upgrade of the Internet infrastructure at their property.

However, there are large heterogeneities in the costs that occur when providing broadband Internet at a specific property. Most importantly, the costs of providing faster Internet is influenced by the distance of a property to the closest main distribution frame (MDF), which is a nodal point in the telecommunication infrastructure. The further the distance to a MDF, the higher the costs of increasing speed. Similar heterogeneities exist in the costs of providing other property features. We therefore introduce a cost parameter  $\beta$  which reflects the cost structure of a seller.

In a simple model, the profit of a seller from renting out a property is determined by the difference between the rent and the costs that occur from providing  $Z$ .

$$\pi = P - C(z_1, z_2, \dots, z_n) \quad (2.4)$$

Given a fixed profit  $\pi$  and cost parameter  $\beta$ , we define the *offer* function of a seller as  $\phi(Z; \pi, \beta)$ . Analogous to the bid function on the buyer side,  $\phi$  represents the prices that sellers are willing to accept for different combinations of  $Z$ .  $\phi_{z_i}$  reflects the marginal rate of substitution between providing any level of  $z_i$  and the increase in the rent that can be obtained on the housing market.

The interactions between sellers and buyers on the housing market will reveal the hedonic price schedule. Wherever the offer and bid functions meet (i.e. have equal gradients), a match between buyer and seller is achieved. The hedonic price schedule is the market clearing envelope of all offer

and bid functions. At every point on the price schedule, a buyer's marginal willingness to pay for  $z_i$  equals the seller's marginal costs incurred by providing one more unit of  $z_i$ . Bid functions vary from buyer to buyer due to different preferences and offer functions vary from seller to seller due to heterogeneities in the cost function.

Figure 2.1 depicts the hedonic price schedule with offer and bid functions for one attribute, namely Internet speed. The functions represent the combinations of Internet speed  $z_{internet}$  and price  $p$  at which buyers and sellers are indifferent. All other attributes  $z_i$  are held constant at their optimal level  $z_i^*$ . For sellers, we therefore obtain offer functions of the form

$$\phi(z_{internet}, z_2^*, \dots, z_n^*; \pi^*, \beta^*) \quad (2.5)$$

where given constant profits  $\pi^*$  and a fixed cost function  $\beta^*$ , the marginal rate of substitution between  $z_{internet}$  and the rent are represented by  $\phi$ . In Figure 2.1,  $\phi^1$  and  $\phi^2$  are two different sellers with different cost functions  $\beta$ . The main reason for the differences in the cost functions is that the property of  $\phi^2$  is closer to the next MDF than the property of  $\phi^1$ . Due to the larger distance to the MDF, it is considerably more costly for  $\phi^1$  to provide high Internet speeds.  $z^{bb}$  denotes the level of  $z$  that we consider as broadband. In practice, properties that are more than 4200 meters away from their MDF can only receive broadband after major infrastructure upgrades. For buyers, the bid functions for Internet speed have the form

$$\theta(z_{internet}, z_2^*, \dots, z_n^*; u^*, I^*) \quad (2.6)$$

where given constant utility  $u^*$  and fixed income  $I^*$ ,  $\theta$  represents the marginal rate of substitution between  $z_{internet}$  and money that can be spend on other goods. The two bid functions  $\theta^1$  and  $\theta^2$  are different due to differences in the preference for Internet speed. For  $\theta^1$ , the marginal willingness to pay for Internet speeds that are higher than  $z^1$  is close to zero.  $\theta^2$  has a preference for higher speeds and will buy  $z^2$ . The tangencies of the two offer and bid curves shown in Figure 2.1, represent two different levels of Internet speed, one being broadband and the other non-broadband, with their respective marginal prices. The matching of all offer and bid curves on the housing market will result in the market clearing hedonic price schedule for Internet speed.

The hedonic price schedule can be used to infer the theoretic welfare effects of changes in broadband availability that are exogenous to buyers and sellers. We will treat the investment decisions of network providers as such exogenous events that are neither anticipated by sellers, nor

by buyers. In the equilibrium depicted in Figure 2.1, seller  $\phi^1$  will provide lower Internet speeds because of his cost function that is different from  $\phi^2$ . For properties that cannot get broadband Internet due to their high distance to the MDF, there are various possibilities to upgrade the network such that broadband becomes available. A first step is usually to connect the MDF and the street cabinet which is closest to the property with fibre optic cable. If a network provider decides to connect a street cabinet with fiber, all properties connected to the cabinet will get higher Internet speeds.

The investment of a network provider will therefore move the offer curve  $\phi^1$  to a position beyond the broadband level  $z^{bb}$ , for example to  $\phi^2$ . At any level of  $p$ , the seller will now provide higher Internet speeds. Since the investments by the network providers are made for every street cabinet individually, we can assume that the upgrade of a single cabinet will not shift the hedonic price schedule due to a small increase in supply of properties with broadband.

On the buyer side, we can expect that consumers will move between properties due to the changes in broadband availability.  $p_2$  will exceed the willingness to pay of buyer  $\theta^1$  which is why he will move to a different property that has a lower level of  $z_{internet}$ . Buyers that were previously not interested in the property of  $\phi^1$  due to their preferences for higher Internet speeds, will now consider to rent this property. If we assume sufficient demand for high Internet speeds, constant preferences and no moving costs, buyers in the housing market will sort such that a new equilibrium is reached in which the price of a property will reflect the new value of  $z_{internet}$  that is possible due to investments by the network provider. In the new equilibrium, utility of consumers is the same as it was before the investment.

In our hedonic model, only the welfare of sellers is affected by the investment of the network provider. The housing price increase from  $p_1$  to  $p_2$  is the gain of sellers, while in the new equilibrium, consumers will pay the same rent to get their preferred level of  $z_{internet}$  at constant utility. The same is true, if the investment into the network is financed by the state and not by the network provider. This finding is common for place-based policies. For example, Kolko (2012) shows that the broadband expansion has led to population and employment growth, but not to increases in wage or a reduction in the unemployment rate. In the end, owners of the immobile resource, i.e. property owners, are the ones that benefit from exogenous infrastructure upgrades, while consumers will move in order to obtain their optimal levels of  $Z$  and  $P$ .

## 2.3 Data

In hedonic regressions with housing prices, three different sources for price data are commonly used.<sup>11</sup> The most widely used data source are government records on house transactions that are obtained from deed transfers. Although this is the most complete data source for house sales, the number of observed housing characteristics is often limited. Data from tax assessments, building permits or mortgage transactions are therefore often used as supplements or substitutes. Another frequently used data source are records from realtor associations. In order to avoid double listings, realtors often share databases in which they list properties they are representing. The obvious drawback of these databases is that they only include properties sold by realtors. There are also concerns about the reliability of the housing prices in these datasets (Parmeter and Pope, 2012). The third commonly used data source are household surveys. Advantages of survey data are that they can usually be assumed to be representative and that they often contain detailed information about housing characteristic. A disadvantage is that they generally contain a lower number of observations than the other two sources.

In this study, we are interested in the hedonic price of broadband Internet. A particularity of broadband availability is that it can in most cases be assumed to significantly improve over time due to investments by network providers and municipalities. In 2005, for example, the average DSL availability rate in German municipalities was 56 percent, in 2013 it reached 94.5 percent.<sup>12</sup> Unlike the value of schools or other more permanent amenities, the value of broadband Internet should therefore not be discounted over the whole life-span of a property. House buyers will assume that sooner or later, higher Internet speeds will become available at the property's location. Since it is unclear for how long broadband availability will be limited, it would be necessary to make assumptions about home buyers' expectations concerning the DSL roll-out when estimating the willingness-of-pay for broadband Internet. In order to avoid making this kind of speculative assumptions, we use property rents data instead of property sales data. Monthly property rents are likely to reflect the current broadband availability better than sales prices.

Besides household surveys, none of the above-mentioned conventional data sources used in hedonic studies provide information on rent prices. We therefore use unique observational data from ImmobilienScout24, Germany's largest online platform for real-estate advertisements. The

---

<sup>11</sup> See Pollakowski (1995) and Parmeter and Pope (2012) for detailed discussions on commonly used sources for housing data.

<sup>12</sup> Data from *Breitbandatlas*.

platform allows potential sellers and landlords to advertise objects that are available for sale or renting. ImmobilienScout24 estimates that about 50% of all real estate transactions in Germany have previously been advertised on their platform Georgi and Barkow (2010). The dataset used in this paper covers the years from 2007 to 2009 and only includes properties that are available for rent.

The original data has monthly frequency. We observe a property in a given month, if it has been advertised at any time during the month. As we are not able to directly observe the price that the property is finally rented out for, we approximate the final rent by the last price the property has been advertised for. We include each property only once, at the time the advertisement is withdrawn. If an object is advertised, then withdrawn from the website and reentered at a later time, we only use the last observation. This leaves us with about 1.38 million advertisements over three years.

Besides the rent price of the property, we observe a variety of object specific characteristics.<sup>13</sup> Most importantly, we observe in which municipality the property is located. Regarding the quality of the property, we observe the living space in square meters, the number of rooms, the year of construction, the type of property,<sup>14</sup> the condition of the property,<sup>15</sup> and whether or not the property has amenities, such as a basement, a garden, an elevator, a balcony, or a build-in-kitchen.

Figure 2.3 shows a screenshot of a typical property advertisement on ImmobilienScout24. Area A shows the title and the full or approximated address of a property. Area B presents photos and area C provides the price and basic characteristics of the property. Note, that there is no specific field where sellers provide information on broadband availability. However, there is a category "TV / Internet / Telefon" with a link to a telecommunication provider that allows to check the availability at the location of the property. There is also information on all extra charges and the total rent price including these charges. In our analysis, we will use the net rent price excluding extra charges. In area D, sellers can provide a more detailed description of the property or provide information on additional amenities. As shown in Figure 2.3, this is the place where some sellers also provide voluntary information on broadband availability. Area E shows the name, address and average rating of the seller. Since many sellers are real estate brokers, it is not uncommon for them to have a large number of reviews from previous buyers. Area F provides a contact form.

---

<sup>13</sup> For a more detailed description of the dataset, see Bauer et al. (2013).

<sup>14</sup> We observe 10 different property types: top floor, loft, maisonette, penthouse, terrace flat, floor apartment, apartment, mezzanine, basement, and other.

<sup>15</sup> We observe 9 different condition categories: first occupation, as new, renovated, in need for renovation, modernized, by arrangement, cares, first occupancy after modernization, and redeveloped.



As we are dealing with information that is manually entered into an user interface, errors in the entry process are a potential issue. To exclude such erroneous entries, we follow a very general approach and exclude the lowest and highest one percent of the observations for the rent price and the living space. For the number of rooms we exclude the top one percent as well as all observations with less than one room. Our results are robust to narrower outlier definitions and also hold when all observations are included.

Table 2.1 provides some descriptive statistics. Columns 2-6 show the minimum, maximum, mean, standard deviation and number of observations for all apartments in our sample. The average monthly rent price amounts to €468.29 and the average property size is 71.69 square meters. Columns 7-15 show statistics for the three subsamples that we use for our instrumental variables estimations. The details on how we obtain each subsample are described in Section 2.4. As the first two subsamples are restricted to west Germany, the average rents for these samples are higher than for the full sample. The third subsample, which is restricted to properties in eastern Germany, has a considerably lower average rent. The number of observations show that each subsample only represents a fraction of the full sample. One reason for this is that all subsamples exclude cities and areas of high population density. In less densely populated areas, the share of rented properties is considerably lower as many people prefer to buy a house or apartment. Also the share of arm-length transactions between family and friends might be higher in rural areas.

Since the actual DSL availability of properties cannot be observed in advertisements on ImmobilienScout24, we use an approximate measure of DSL availability in the property's municipality. In Germany, municipalities are the smallest territorial divisions. Between 2007 and 2009, there were more than 12,000 German municipalities with a median population of about 2600 inhabitants and a median area of about 25 square kilometers. For every municipality, we observe broadband availability in 2007, 2008 and 2009 from the German *Breitbandatlas*. The *Breitbandatlas* collects information on the availability of different broadband technologies from all telecommunication providers in Germany. It then estimates the overall availability for every technology on the municipality level. In this paper, we use information on the availability of DSL, which is by far the most common fixed-line broadband access technology in Germany. In some specifications, we also use the availability of mobile broadband technology as a control. Since broadband availability does usually not vary on the level of a single property but mostly affects a whole neighborhood, measuring availability on a more aggregate level is a good approximation. One possible interpretation of the broadband availability rate for a single property is the probability of the property to have broadband available.

## 2.4 Empirical Strategy

In this Section, we outline our strategy to estimate the hedonic price for broadband Internet. We start with a general description of our estimation model. Since our identification is based on the historical roots of broadband Internet infrastructure, we briefly describe the history of broadband in Germany and explain a few technical particularities that are important for our instrumental variables. Our estimation strategy generally follows the strategy of Falck, Gold and Heblich (2014).

### 2.4.1 Estimation of the Hedonic Price for Broadband

Estimating the marginal willingness-to-pay for a property amenity using a conventional hedonic pricing model is econometrically not very demanding. We regress housing prices on all observable property amenities. The  $\beta$ s associated with each amenity can then be interpreted as the marginal willingness-to-pay in the hedonic equilibrium described in Section 2.2.

Therefore, we estimate:

$$\log(\text{rent}_{im}) = \beta_0 + \beta_1 \text{DSL}_m + \beta_2 X_i + \beta_3 Y_m + \varepsilon_{im} \quad (2.7)$$

where  $\text{rent}_{im}$  is the monthly rent of property  $i$  in municipality  $m$  in euros.  $\text{DSL}_m$  is a continuous variable indicating the DSL availability at the municipality level. Vector  $X_i$  includes a set of property characteristics, such as the number of rooms, the age of the property and if the property has a garden, terrace or basement. Vector  $Y_i$  is a set of regional characteristics, such as the population size, tax revenues and average age in the municipality.

The problem of this estimation is that amenities, such as DSL, are typically not randomly distributed across properties. Areas with higher spending power are likely to also have higher demand for broadband Internet. Network providers will react to this demand and foster broadband availability in these areas. As spending power is likely to be also correlated with a number of other amenities, such as better schools or more cultural offerings, it becomes difficult to identify the effect of broadband availability. Thus, estimations of the form presented in Equation 2.7 will presumably suffer from omitted variable bias.

In order to avoid this potential bias, modern hedonic pricing studies exploit exogenous variation and apply difference-in-difference, regression discontinuity, or instrumental variable approaches.<sup>16</sup> These identification strategies are especially popular for estimating the value of environmental amenities, such as hazardous waste (e.g., Gayer, Hamilton and Viscusi, 2000; Bui and Mayer, 2003; Greenstone and Gallagher, 2008), air quality (e.g., Chay and Greenstone, 2005), noise (e.g., Pope, 2008*a*) and other health risks (e.g., Davis, 2004). Quasi-experimental settings have also been used to estimate the hedonic price of school quality (e.g., Black, 1999; Gibbons and Machin, 2001; Figlio and Lucas, 2004), disaster risks (e.g., Hallstrom and Smith, 2005; Pope, 2008*b*) and crime (e.g., Linden, 2008; Pope, 2008*c*).

In this paper, we apply an instrumental variable approach that exploits exogenous variation in DSL availability. The variation that we are using stems from the preexisting voice telephony network which was designed before the invention of broadband Internet. All instrumental variables in our analysis are related to the distance or other characteristics of nodal points in the voice telephony network, called MDFs. Our first stage estimations are therefore of the form:

$$DSL_i = \beta_0 + \beta_1 MDF_i + \beta_2 X_i + \beta_3 Z_i + \varepsilon_i \quad (2.8)$$

where  $MDF_i$  is the respective instrumental variable that is based on distances to and characteristics of MDFs. In the following subsections, we describe the construction of each instrumental variable in detail.

## 2.4.2 Historical Features of Broadband Infrastructure in Germany

In order to understand the main determinants of broadband availability and the sources of exogenous variation we are going to exploit in this paper, it is useful to take a look at the telephone network in Germany before the arrival of the Internet.<sup>17</sup> The first German telephone networks were installed at the end of the nineteenth century. Until 1908, all calls were manually connected by human operators at telephone exchanges. With the introduction of electromechanical exchanges, calls within the same local line network could be automatically routed. For calls outside the local line network, human operators were still necessary. These operators were located at the telephone exchanges, which are the points where the dedicated copper wires for every household come together. The

<sup>16</sup> For a detailed discussion on quasi-experiments in hedonic models see Parmeter and Pope (2012).

<sup>17</sup> Since we use data from Germany for our analysis, we will focus on the German telecommunication infrastructure. However, the history and basic layout of the telephone network is very similar in other countries.

copper wires between every household and the telephone exchange are often referred to as the "local loop" or the "last mile" of the telephone network.

Starting in 1923, electromechanical routing was not only used to connect households within a local loop, but also to connect different local loops. The manual telephone switchboards were step by step replaced by so-called Main Distribution Frames (MDF) which are used for the more permanent wiring of the dedicated copper cables with other telecommunication equipment. Despite early progress, it took very long until the whole telephone network was routed completely automatically. In West Germany, human operators at some telephone exchanges were in service until 1972. In a few rural East German areas, human operators were in service until the late 1980's. Today, human operators and manual exchanges are history, but the design of the telecommunication network, especially the local loop, still goes back to the origins of voice telephony.

For voice telephony, the distance between a household and the local exchange did not play a critical role. In general, signal loss in copper cables is dependent on the frequency of the signal. The higher the frequency, the higher the signal loss. Since the analog transmission of voice uses a very low frequency band, the signal loss in the old telephone network was relatively low, even if copper cables were several kilometers long. When the Internet became popular in the 1990's, analog modems were used for data transmission over the telephone network. These modems converted the digital data from computers into analog signals that used the same frequencies as voice telephony. Every telephone line could therefore also be used for Internet access.

The increasing demand for higher Internet speeds led to the development of the Digital Subscriber Line (DSL). The idea behind DSL is to use frequency bands that are higher than those of voice telephony to transmit data digitally through copper cables. Using higher frequencies increases the bandwidth for data transmission and therefore increases Internet speed. The problem is that higher frequencies also suffer from higher signal loss. Thus, the length of the copper wire became suddenly important. Beyond a certain length, the signal loss of DSL becomes so high that the line cannot be used for DSL data transmission. The shorter the distance between the MDF at the local exchange and the household, the larger the frequency band that can be used for data transmission. For this reason, the decades-old telephony network which was never built for the transmission of signals with higher frequencies than voice, became the limiting factor for Internet speed.

Figure 2.2 shows the local loop, sometimes called the "last mile", of the telephone network. For the historical reasons described above, each household has a dedicated copper wire which terminates

at the MDF. Street cabinets between the household and the MDF merely bundle a set of copper wires into a larger jacket. This network design made sense in the time when the copper wires were only used for voice transmission and the local exchange was the location where human operators connected telephone calls. If a copper wire network was planned today, one would probably choose to terminate the dedicated copper cables in the street cabinets and place the MDF and corresponding telecommunication equipment there. This way, the maximum length of a copper wire between household and MDF would be drastically reduced to a few hundred meters.

The historical "design flaw" of the telecommunication network can be used as exogenous variation for identifying the effect of DSL availability. In Germany, DSL is by far the most common fixed-line broadband Internet access technology. Cable networks, which play an important role as Internet providers in other countries, are not very dominant in Germany. Figure A2.3 shows that in 2007, 94 percent of all broadband subscriptions were DSL. Although the share of other broadband access technologies increased over time, even in 2009, the latest year of our analysis, 9 out of 10 broadband subscriptions were DSL. We will therefore treat the terms "broadband" and "DSL" interchangeably in this paper.

### 2.4.3 Exploiting Variation in the Distance Between Properties and MDFs

In rural areas, one MDF often covers more than one municipality. During the time when human operators manually connected telephone calls, it was simply not economical to have one exchange for every small municipality. Even today, about 25 percent of all municipalities have less than 1000 inhabitants. Decades ago, when the layout of the voice telephony network was determined, Germany had even more municipalities with an even lower average population. Therefore, many municipalities were connected to the same MDF and consequently shared the same prefix. Today, there are around 5,200 prefix areas in Germany. Since in cities, one prefix area is often served by several MDFs, our sample includes 7,971 MDFs. In rural areas, a prefix area usually has exactly one MDF.

About 58 percent of German municipalities do not have an own MDF. By restricting our sample to properties in municipalities without MDF, we implicitly impose the condition that all remaining properties are located in less densely populated areas. The remaining municipalities share their MDF with about three other municipalities, on average.

Figure 2.4 shows the relationship between the average broadband penetration in municipalities and the distance from municipality centroids to the MDF for 2007-2009. The first graph shows

that municipalities that are more than 4000 meter away from their MDF have significantly lower broadband penetration rates. There is only a very little trend below the threshold and a very significant drop above the threshold. Theoretically, this threshold should be at about 4200 meters, as this corresponds to the maximum attenuation of 55dB in average copper wires. When the attenuation is higher than 55dB, the signal loss is too high for DSL. Depending on the cable route, the cable diameter and the cable quality, the attenuation can be higher or lower, such that the threshold is fuzzy.

The second graph in Figure 2.4 shows the same relationship in a sample of municipalities without an own MDF. This restriction basically cuts away the lower part of the distance distribution. Most of the remaining municipalities have distances between 2000 and 6000 meters. For our final estimation sample we make two additional restrictions. First, we exclude all eastern German municipalities. After reunification, the telephone network in eastern Germany was subject to large reorganizations leading to a couple of particularities which we will discuss later. The second restriction is that we exclude all municipalities above the threshold that have another MDF nearby. We will discuss this particularity in detail below. The resulting estimation sample is depicted in the third graph of Figure 2.4. Henceforth, we will refer to this sample as the distance sample.

Our first instrument for the estimations with property data is a dummy variable that takes on the value 1 for municipalities with distances above the 4200 meters threshold. Figure 2.4 clearly shows the relevance of this instrument. The relatively sharp drop above the threshold and the very little trend below the threshold are evidence for the exogeneity of the instrument. This argument is strengthened by repeating the estimations in a sample that is restricted to municipalities around the threshold. By including MDF fixed effects, we only compare properties in neighboring municipalities that do not have an own MDF. For example, one property is located in municipality A that does not have an own MDF. Another property is located in the neighboring municipality B which has the same prefix than municipality A and also does not have an own MDF. Since the distance between municipality A and the MDF is lower than 4200 meters, the property in municipality A can get DSL. The distance between municipality B and the MDF is higher than 4200 meters which means that the property in municipality B cannot get DSL.

One way to interpret our first instrument is that it exploits differences in the roll-out costs of DSL. For network providers, the lowest roll-out costs arise in municipalities that are close to a MDF. In these municipalities, the only necessary investment is the new hardware at the local exchange. For areas that are more than 4200 meters away from a MDF, roll-out costs are considerably higher.

Usually, the street cabinets depicted in Figure 2.2, have to be connected with fibre wires in order to reduce the length of the copper wires to the households. Since in Germany almost all cables are buried, the costs of connecting a street cabinet with fibre are very large. Although the costs depend on many parameters, they are often in the ballpark of about €80,000 per kilometer. Thus, the roll-out of DSL to municipalities above the threshold is considerably more expensive than the roll-out to areas below the threshold.

#### **2.4.4 Exploiting Variation in the Roll-out Costs for Properties that are Far Away from their MDF**

For some municipalities above the threshold, roll-out costs are slightly lower than for others. As described in Section 2.4.2, political boundaries often played an important role when it was decided which household should be connected to which MDF. When one municipality belonged to another municipality politically, both municipalities were expected to have the same prefix. Therefore, both municipalities had to be connected to the same MDF. In some cases, this MDF was not the closest MDF for the municipality. In the voice telephony world, the distance to the MDF was not a critical criterium and thus it was not perceived crucial to connect municipalities to their closest MDF. In the DSL world, municipalities that are not connected to their closest MDF can often not get broadband. However, if another MDF is close by, the roll-out costs for these municipalities are still lower than for municipalities that do not have a closer MDF. The reason is that the street cabinets can be connected to the closer MDF.

Our second instrumental variable exploits the differences in roll-out costs among the municipalities above the threshold of 4200 meters. The binary instrument takes on the value 1 if a municipality does not have any MDF closer to the one it actually is connected to. For these municipalities, the roll-out costs are even higher than for municipalities that are above the threshold but can be connected to a closer MDF. We use this instrument in a sample of municipalities that do not have an own MDF and are connected to a MDF that is above the threshold. Henceforth, we will refer to this sample as the above-threshold sample.

While we excluded municipalities above the threshold with a closer MDF nearby in the distance sample, these municipalities are key to our estimations in the above-threshold sample. Indirectly, Figure 2.4 shows that when municipalities with a closer MDF nearby are excluded, the broadband penetration is more stable over time. In the second graph, the vertical distance of the points above the threshold is relatively large. This shows that between 2007-2009 many of the municipalities

increased DSL penetration. In the third graph, which excludes municipalities with a closer MDF, the development over time has been less dramatic. Especially between 2007 and 2008, the DSL penetration improved only very little. One reason for this difference between Graph 3 and 4 is that Graph 3 includes municipalities with lower roll-out costs due to a closer MDF.

Note that in the above-threshold sample, it would be very restrictive to include MDF fixed effects. Among those municipalities without an own MDF, one MDF is on average shared by four municipalities. In most cases, only one or two of these municipalities are above the threshold. By including MDF fixed effects, we would only identify over those municipalities that have a closer MDF than the one they are connected to and have at least one neighboring municipality that is also above the threshold but does not have a closer MDF than the one it is connected to. In order to be a little bit less restrictive, we include county fixed effects instead of MDF fixed effects when we use the second instrument. We therefore compare all municipalities without an own MDF, above the 4200 meters threshold, in the same county. We also include the distance between the municipality and the closest urban center as an additional control that should capture if a municipality only has a closer MDF because it is closer to another urban center.

#### **2.4.5 Exploiting Variation from a Technological "Mistake" in Eastern Germany**

A third source of exogenous variation in DSL availability comes from a historical "mistake" that was made in the East German telecommunication infrastructure after reunification in 1989. Since households in East Germany often had very poor telephone lines, the state-owned telecommunication monopolist *Deutsche Bundespost* decided to replace parts of the eastern German infrastructure with the then state-of-the-art Optical Access Line (OPAL) technology. The selection of OPAL areas was as good as random.<sup>18</sup> At that time, the fibre-based OPAL technology was considered to be superior over traditional copper lines. Before the arrival of DSL technology, the OPAL fibre wires had indeed several advantages compared to the older voice telephony network. However, it later turned out that OPAL is not compatible with DSL which relies on traditional copper wires.

Due to the incompatibility with DSL, 213 affected OPAL MDFs serving about 11 percent of the eastern German households could not get broadband Internet when DSL became available. In order to provide these households with DSL, network provider either have to replace the fibre wires

---

<sup>18</sup> For a more detailed discussion about the exogeneity of OPAL, see Gebhardt (2010) and Bauernschuster, Falck and Woessmann (2014).



with copper or need to install expensive hardware at the MDF or street cabinets. We exploit this variation in DSL availability by constructing a binary instrument that becomes unity if a household is connected to an OPAL MDF. We restrict our sample to eastern German households that are closer than 4200 meters to their MDF. Only for municipalities below this threshold, DSL would be available if the MDF was not of the OPAL type. We also exclude municipalities with more than one MDF in order to ensure that the property in question is connected to either an OPAL or a non-OPAL MDF. Henceforth, we will refer to this third sample as the OPAL sample.

## 2.5 Results

In this section, we present and discuss results of our hedonic regressions of property rents on DSL availability. We first describe the general results, then discuss the interpretation of the estimate sizes and finally compare them to estimates from previous studies.

### 2.5.1 The Effect of DSL Availability on Property Prices

Table 2.2 shows results from estimations in our distance sample of properties. All properties are located in municipalities that do not have an own MDF, but are connected to the closest MDF of another municipality. The coefficients in the first line show the effect of increasing broadband availability by 10 percent. As shown in Table 2.1, the standard deviation of DSL availability in this sample is 8 percent. Therefore, the observed estimates correspond to an increase in DSL availability that is larger than one standard deviation. Column 1 and 2 show OLS estimates from two different specifications. In the first column, an increase of DSL availability by 10 percent leads to a modest increase in house prices by 0.57 percent. Controlling for mobile broadband availability in the second column hardly changes the result. All estimations include a set of property characteristics, such as building age or the availability of a balcony. Additionally, we control for a wide set of municipality characteristics and include MDF fixed effects. The latter ensures that we only compare properties in municipalities that are connected to the same MDF.

Column 3 and 4 of Table 2.2 show the reduced form effect of the first instrumental variable on DSL availability. The instrument is a dummy variable that takes on the value 1 if the distance to the MDF is larger than 4,200 meters. Being above this threshold has a negative effect of 1.47 percent on rent prices. It is hardly effected by controlling for wireless broadband availability. Column 5 and 6 show estimates in which DSL availability is instrumented with the first instrumental variable.

The IV estimates are about 7 times larger than the OLS estimates and indicate that an 10 percent increase in DSL availability leads to about 4 percent higher rent prices. Possible explanations for the difference in the estimate size are discussed in Section 2.5.2.

Table 2.3 shows results for our above-threshold sample which only includes properties in municipalities that are further than 4,200 meters away from their MDF. This sample restriction reduces the observations to less than a third of the distance sample. The first columns shows that the OLS estimate is considerably larger than in the distance sample. This could indicate that for more "disadvantaged" municipalities with lower DSL availability, the effect of broadband availability on rent prices is higher. In this sample, also the estimate of wireless broadband availability is larger than in the distance sample. After controlling for mobile broadband, the estimate of DSL becomes considerably smaller. A possible explanation would be that the wireless network is not independent from the fixed-line network. Every cellular network antenna needs to be connected to some broadband Internet line,<sup>19</sup> favorably to a high-bandwidth fibre wire connected to the closest MDF. In our above-threshold sample, the closest MDF is at least 4,200 meters away which also affects mobile broadband availability negatively. Thus, for properties in disadvantaged municipalities with little DSL availability, the correlation between DSL and mobile broadband availability is higher.

Column 3 of Table 2.3 presents the reduced form estimate for our second instrument. The instrumental dummy variable becomes 1 if municipalities do not have any other MDF nearby that is closer than 4,200 meters and could therefore reduce upgrade costs for municipalities in the above-threshold sample. The coefficient has the expected negative sign. After controlling for mobile broadband availability, the coefficient becomes smaller in size and is not statistically significant anymore. Again, this could be explained by the dependence of the mobile network on the fixed-line network described above. Columns 5 and 6 show the IV estimates for our second instrumental variable. The estimates are similar to those of the first IV in the distance sample.

Table 2.4 shows results for our OPAL sample which only includes properties in eastern German municipalities that are less than 4,200 meters away from their MDF. Surprisingly, in this sample the OLS estimate of DSL availability is basically zero. The reduced form coefficients in Column 3 and 4 are negative, as expected. The IV estimates are considerably smaller than the estimates in the previous two samples, but are significantly different from zero. A 10 percent increase in broadband availability has a positive effect of 1.5 percent on rent prices in the OPAL sample.

---

<sup>19</sup> In some cases, cellular network antennas are connected to other cellular antennas by point-to-point radio links.

There are two major differences between the first two samples and the OPAL sample. First, the OPAL sample only includes eastern German properties, and second, the OPAL sample only includes properties closer than 4,200 meters from the MDF. In order to understand which of the two differences drives the differences in estimate size, we can compare OLS estimates for eastern and western Germany. Unfortunately, the first two instrumental variables cannot be applied in eastern Germany due to the differences in the history of the telecommunication infrastructure. Since many parts of the voice telephony network had to be redesigned and replaced after reunification, the assumption that our first two instruments rely on do not hold in eastern Germany. Similarly, the historical mistake of rolling out OPAL technology can only be exploited in eastern Germany. Therefore, we cannot directly compare IV estimates for eastern and western Germany.

Table 2.5 shows OLS estimates for different samples in eastern and western Germany. In the full sample, there is a positive and significant association between DSL availability and rent prices for properties in west Germany. The size of the coefficient is comparable to the size of the respective coefficient in the distance sample. However, this positive association cannot be observed for east German properties. Also in the distance sample, we do not find a positive association between DSL availability and rent prices for eastern German properties. In the above-threshold sample, the OLS estimate is positive and significant for both west and east German properties. In the OPAL sample, we neither find a positive association for east, nor for west German properties.

The estimates in Table 2.5 indicate that on the one hand, the difference in results between the first two samples and the OPAL sample go back to differences in effects between eastern and western Germany. On the other hand, they seem to go back to differences in effects for properties below and above the 4,200 meter threshold. In east Germany, we neither observe a positive association for the average property, nor for properties below the threshold. But we do find a positive association for east German properties above the threshold. In west Germany, we observe a positive association for both the average property and for properties above the threshold. For properties below the threshold we cannot observe a positive association. The differences between east and west can probably be explained by the particularities of the east German housing and labor market. After reunification, east Germany underwent major transitions that affect the housing and labor market until today. Possible explanations for the differences between properties above and below the threshold are discussed in the next section.

## 2.5.2 Interpretation of the Estimates Size

Our IV estimates in all three samples are considerably larger than the respective OLS estimates. There are several possible explanations for these differences in coefficient size. In this Section, we discuss possible LATE interpretations of the IV coefficients as well as a few particularities of the broadband availability data that could lead to an underestimation of the OLS coefficients.

The OLS estimates in our regressions can be interpreted as the average association between DSL and rent prices in the respective sample. This association seems to be stronger for properties above the 4,200 meter threshold and weaker for properties below the threshold. In each sample, the OLS estimates compare properties along the whole distribution of distances and corresponding DSL availability rates. This distinguishes them from our instrumental variable estimates. Since not all properties in our samples have the same likelihood to respond to our instruments, we can only identify the local average treatment effect (LATE) for those properties that are influenced by changes in the instrumental variables.<sup>20</sup>

Our instruments only affect the treatment status of municipalities that are at the lower end of the DSL distribution. In the case of our first instrument, being above the 4,200 meters threshold reduces the chances of getting any DSL at all. For our second instrument, having no other MDF close by also reduces the chances to get even low-speed DSL access. And for our third instrument, being connected to an OPAL MDF means that the affected properties will not get Internet speeds faster than ISDN. Municipalities with high DSL availability are affected to a much smaller extent by the instruments. Among these municipalities at the top of the DSL distribution, the major variation lies in the availability of different DSL speeds. Municipalities with higher availability of basic DSL, also have a higher share of high-speed DSL connections.<sup>21</sup> In other words, our instruments are most relevant at the *extensive* margin of either having DSL or not. The OLS coefficients, however, also measure the effect along the *intensive* margin of having slower or faster DSL access. One possible interpretation of the differences between our OLS and IV estimates is therefore that the extensive margin is of higher importance than the intensive margin.<sup>22</sup>

<sup>20</sup> For a formal discussion of local average treatment effects, see (Imbens and Angrist, 1994).

<sup>21</sup> This relationship mainly goes back to the distance to the MDF and the population density within municipalities. In our data from 2007-2009, we do not observe the availability of high-speed DSL access. For 2010, the *Breitbandatlas* reports the availability of different DSL speeds on the state level. Figure A2.1 shows that states with a shorter weighted average distance between municipality centroids and MDFs also have a higher share of 16 Mbps over 1 Mbps connections available.

<sup>22</sup> This interpretation is supported by the notion that Internet use changes the most at the extensive margin. Table A2.2 presents results from estimations with SOEP data on the individual level. It shows how the Internet use intensity on a scale from "never" to "daily" changes with increasing distance of the household to the MDF. Compared to the base category of households that are closer than 1,000 meters to the MDF, the most significant drop in use intensity can

The fact that the OLS estimates in the above-threshold sample are considerably larger than the OLS estimates in the distance sample supports our proposed LATE interpretation. The above-threshold sample only contains properties in municipalities that already have low chances of broadband availability due to their large distance to the MDF. Among these properties, the question is not so much about DSL speed, but rather if they get any DSL at all. Note that the IV estimates remain relatively constant between the distance and the above-threshold samples. This indicates that the instruments affect similar properties in both samples, while the OLS estimates measure average effects in different parts of the DSL distribution.

There are also reasons why OLS might underestimate the effect of DSL availability on rent prices. One possible reason is that the minimum broadband speed definition used in the *Breitbandatlas* is very low. In 2007, every connection that is faster than ISDN access (0.125 Mbps) is considered to be broadband and in 2008-2009 the minimum broadband definition is 0.375 Mbps. However, most consumers would probably only consider speeds above 1-2 Mbps as broadband. In 2008, only 6 percent of all broadband subscriptions in Germany were slower than 2 Mbps (Bundesnetzagentur, 2008). For many consumers, the availability of broadband access as defined in the *Breitbandatlas*, might therefore overstate the availability of broadband access with typical minimum DSL speeds. This overassessment would not necessarily be constant across regions. Due to the correlation between broadband availability and speed, areas with high availability are also likely to have a higher share of average broadband speeds above 1 Mbps and areas with low availability are likely to have a higher share of broadband speeds below 1 Mbps. As a result, the OLS coefficients of broadband availability might underestimate the effect of the availability of broadband with conventional DSL speeds.

Another concern with the *Breitbandatlas* data is that there might be measurement error due to overreporting. In order to calculate the availability for every German municipality, the *Breitbandatlas* collects data from telecommunication providers that indicate their broadband coverage in a given area. One could argue that providers usually have incentives to over- rather than understate their coverage. Before the latest amendment of the German telecommunication law in 2012, there was a lively debate if the amendment should include a Universal Service provision for broadband Internet. This provision would have meant that the incumbent is obliged to provide broadband to all areas where broadband is not yet available. Both the incumbent and its competitors were opposing this provision for obvious reasons. Thus, providers had an interest in understating the problem of

---

per observed for individuals that are very far away from the MDF. If the value that individuals put on broadband Internet is related to its use intensity, than the results in Table A2.2 could support the hypothesis that broadband is more valuable at the extensive margin.

insufficient broadband availability. If the data from *Breitbandatlas* systematically overestimated broadband availability in areas with low coverage, this would lead us to underestimate the availability effect in OLS.

### 2.5.3 Comparison of our Estimates with Estimates from Previous Studies

Our largest estimate, the IV estimate in the above-threshold sample, indicates that an increase of broadband availability by 10 percent leads to 4.7 percent higher rent prices. Given the average rent of €508.79 in this sample, the monetary value would amount to €23.9 per month. For properties in eastern Germany and properties at the upper end of the DSL distribution, the effects are likely to be considerably smaller. In order to put our results into perspective, it is useful to compare them to the estimates of previous studies. As shown in Section 2.1, there is a wide variety of methods that are used to estimate the value of broadband Internet. For example, Greenstein and McDevitt (2012) estimate the consumer surplus of broadband Internet for all OECD countries and find that in Germany, the quality-adjusted broadband bonus per subscriber amounted to \$ 496.12 in 2010. This corresponds to a consumer surplus of \$ 41.34 per month. Using a conjoint approach, Rosston, Savage and Waldman (2010) find that U.S. consumers are willing to pay \$ 59 per month for basic Internet access and \$ 98 per month for a premium access with increased speed and reliability. A similar methodology is applied by IAB Europe and McKinsey (2010) who additionally account for the costs of paid services and for the disturbance from advertisements, to find that the consumer surplus of broadband Internet per household amounts to €38 per month. Varian (2011) uses estimates from Chen, Jeon and Kim (2014) and data from Google to calculate that the average savings to consumers generated by search engines amount to \$ 500 per year or \$ 41.66 per month. Valuing the Internet by the time people spend using it, Goolsbee and Klenow (2006) find that the Internet's consumer gains exceed \$ 3,000 per year or \$ 250 per month.

In the study that is most closely related to ours, Ahlfeldt, Koutroumpis and Valletti (2014) use a hedonic model with house sales data and estimate that a 10 percent increase in Internet speed leads to an 1.2 percent increase in house sales prices. Given the mean property price of their 2005 transaction prices sample, this corresponds to \$4,296. If a house buyer assumes that Internet speed will be limited for the next ten years, this would translate into a price premium of \$429 per year or \$35.8 per month.

Given the estimates from previous studies, a monthly rent premium of about €23 per month for a 10 percentage points higher broadband availability seems plausible. However, we should

keep in mind that the samples we use to estimate our effects are rural regions and the LATE of our IV estimations is for municipalities that are most disadvantaged in terms of DSL availability. Most other studies do not have comparable sample restrictions and calculate values for the average consumer.

## 2.6 Validation of Results with Survey Data from SOEP

The data from ImmobilienScout24 has one important drawback. Since the properties are untenanted and landlords are not obliged to provide information on broadband availability, we do not observe directly if broadband is available at the property's location or not. In order to validate our results from Section 3.2, we use survey data from the German Socio-Economic Panel (SOEP) that have the advantage of providing information on whether a household has DSL or not.

### 2.6.1 Description of SOEP Data

The SOEP is a representative annual household survey with more than 11,000 households, of which less than 50 percent live in rented properties.<sup>23</sup> The 2008 SOEP wave is the first to include a question on the Internet access type. Specifically, households are asked if they have an "Internet connection without DSL" or a "DSL connection". Note that this question captures the actual DSL takeup by households, as opposed to the availability of DSL at a property's location. In 2008, about 51 percent of all households indicated to have a DSL connection. Among households living in rented properties, DSL penetration is slightly lower. Since we want to ensure comparability with the ImmobilienScout24 estimates and neither the 2007, nor the 2009 SOEP waves include information on DSL, we limit our analysis to the 2008 wave.

SOEP households provide detailed information on the characteristics of their rented property. For example, we know the size, number of rooms, general condition and approximated age of the property. The SOEP periodically asks households for the distance of their property to amenities, such as schools, restaurants, grocery stores and recreation areas.<sup>24</sup> Unfortunately, the 2008 wave does not contain information on the distance to amenities. We therefore merge this information

---

<sup>23</sup> Households that live in a property which is owned by a member of the household are disregarded for our analysis. The same is true for subtenants and residents of nursing homes, student dormitories or other institutional housing offers. We also exclude about 8 percent of households in rented properties that indicate that they only pay a reduced rent.

<sup>24</sup> Respondents answer if they can reach the respective amenity within 10 minutes, in 10-20 minutes, in more than 20 minutes or if the amenity cannot be reached by foot.

from the 2004 and 2009 waves to the 2008 wave. For households that are either missing or have moved between 2004 and 2008 and also between 2008 and 2009, we do not observe the data on distances to amenities.

In order to apply a similar, but due to data limitations less elaborate identification strategy as the one described in Section 2.4, we need to know the distance between each household and the MDF that the household is connected to. In order to comply with data protection provisions, SOEP generally only allows spatial analysis on the state level.<sup>25</sup> However, we were granted *in situ* access to the geocoded household data at the German Institute for Economic Research, DIW Berlin, which administers the SOEP. Using the same spatial data on the German telecommunication infrastructure as in the rest of this paper, we are able to calculate the exact distances between households and their MDF in meters.

Table A2.1 provides descriptive statistics of the SOEP variables for the full sample and for the west and east German samples separately. The last Column shows the difference in means between west and east Germany. The second row of Table A2.1 shows that the average rent in the full sample amounts to €463.22. This is surprisingly similar to the average rent of €468.29 in the ImmobilienScout24 sample. For households in western Germany, the rent is considerably higher than for households in eastern Germany.

## 2.6.2 The Value of Broadband for SOEP Households

Table 2.6 presents results of hedonic regressions in the sample of SOEP households. Column 1 and 2 show estimates of OLS regressions with west German households. The coefficient in Column 1 shows that having DSL is associated with 4.7 percent higher rents. When we exclude households that are located in the city center, the association becomes weaker. Rows 2-4 show coefficients of other amenities for a comparison of the estimate size. The DSL estimates are slightly larger than the estimates of an improvement in a property's condition by one category (on a scale from 1-4) and smaller than the estimates of having a terrace.

Due to the limited number of observations, we are not able to use the instrumental variables presented in Section 2.4 with our SOEP data. However, we can use the continuous distance between a household and its MDF as an instrument. A similar approach is followed by Czernich (2012), Belo, Ferreira and Telang (2013) and Ahlfeldt, Koutroumpis and Valletti (2014). When we use the distance to instrument DSL adoption by the household, we make the strong assumption that

<sup>25</sup> A remote access system allows to go further down on the county level.



the distance only has an effect on house prices through DSL availability. In order to make this assumption more plausible, we include a large set of controls for the distances to other amenities and show a specification in which we exclude all households that are located in city centers. Figure A2.2 underlines the relevance of the distance instrument by showing the share of households with DSL in eight different distance bins. There is a clear drop in DSL adoption for households that are further than 3,750 meters away from their MDF.

Columns 3-4 present the results from estimations in which we instrument DSL adoption with the distance to the MDF. The coefficients are about 3 times larger than the OLS coefficients, but not statistically significant. Columns 5-6 show OLS estimates for households in eastern Germany. Consistent to what we found in the ImmobilienScout24 sample, there is no effect of DSL on property prices in eastern Germany.

Generally, the results from estimations with survey data from SOEP are in line with the results from estimations with ImmobilienScout24 data. The OLS estimates for having DSL are larger than the respective OLS estimates for a 10 percent increase in local DSL availability. This is not surprising, as the actual adoption of DSL is a much more direct measure of DSL availability than the average availability on the municipality level. On the other hand, the adoption is also more likely to be correlated with other unobserved household characteristics. The IV estimates in the SOEP sample are even larger than the respective OLS estimates, although we interpret them with caution due to their large standard errors and potential threats to the exogeneity of the instrument. The IV estimates in the ImmobilienScout24 sample are of similar size as the OLS estimates in the SOEP estimations, but reflect a local average treatment effect that is not directly comparable to the SOEP estimates.

## 2.7 Conclusion

Broadband Internet contributes to consumer surplus in many ways. However, the value of broadband to consumers is difficult to measure with conventional methods, since many services that are offered online do not have a price. Using data on property rent prices in a hedonic model, we estimate that the value of broadband can be as high as €23 per month. This represents the willingness-top-pay above the actual broadband subscription price. While Internet Service Providers are able to extract the price of broadband subscriptions, they are not able to extract the surplus that we estimate in this paper due to competition on the broadband market.

A hedonic model of the broadband market shows that as long as properties are scarce and consumers are mobile, landlords extract the surplus from broadband Internet. When telecommunication providers invest into the infrastructure and thereby make broadband Internet available at properties that did not have broadband before, landlords are able to extract higher rents for these properties. Consumers will sort themselves into properties according to their preferences. After the infrastructure upgrade, the utility of consumers will therefore be as high as before the upgrade.

Although landlords are able to extract parts of the surplus from broadband availability, they often decide not to invest into the infrastructure themselves or in cooperation with telecommunication providers. One reason for this could be a freeriding problem. The most common way to provide DSL to properties without broadband is to connect the street cabinet of the property with fibre cable. When a street cabinet is connected with fibre, in most cases all properties connected to the street cabinet are able to get broadband Internet. As a solution to this problem, landlords of properties that are connected to the same street cabinet would need to coordinate and pay the telecommunication provider for the investment. As this coordination seems to fail in practice, policy makers could choose to subsidize the infrastructure (which they actually do) and recoup the costs through property taxes.

The costs of connecting a single property with broadband Internet vary a lot. In a recent study for the German Ministry for Economic Affairs and Energy, TÜVRheinland (2013) estimate the costs of providing all households in Germany with high-speed broadband Internet (at least 50 Mbps). These costs are likely to be a lot higher than the costs of providing basic broadband speed. TÜVRheinland find that for connecting the first 75 percent of households with high-speed broadband, the average cost per household is about €660. For the households between the 75th and 95th percentile, the costs are €810. For households above the 95th percentile, costs are €3.850 per household. This shows the high heterogeneity in costs that mainly goes back to the distance between a property and the closest MDF. Given our estimates, landlords can extract up to €276 per year for broadband availability. This means that within 7 years, landlords could extract 50 percent of the highest costs for the roll-out of high-speed broadband with 50 Mbps.

## Figures and Tables

Figure 2.1  
The equilibrium hedonic pricing schedule with bid and offer functions




Figure 2.2  
Representation of the "local loop" between households and MDF

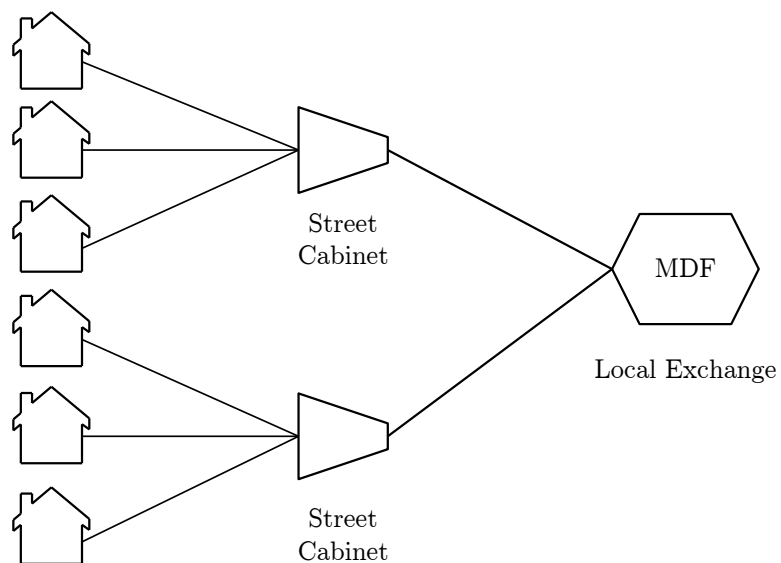


Figure 2.3  
Screenshot of a property advertisement on ImmobilienScout24.de

Scout24 | AutoScout24 | FinanceScout24 | FriendScout24 | JobScout24 | TravelScout24 powered by

**IMMOBILIEN SCOUT 24** Rein ins neue Leben. Willkommen! Anmelden oder [neu registrieren.](#)

**Wohnen** | Gewerbe | Anbieten | Eigentümer



**Suchen** | Suchanzeigen | Baufinanzierung | Hausbau | Wohnideen | Umzug | Markt & Preise | Branchenbuch | Merkzettel | **Mein Konto**

Sie sind hier: [Wohnung mieten](#) > [Sachsen-Anhalt](#) > [Magdeburg](#) > [Sudenburg](#) > [Exposé](#)

## DG-Wohnung in Sudenburg **\*\*DSL verfügbar\*\*** A

39112 Magdeburg, Sudenburg  Umzugskosten vergleichen  
[Karte ansehen](#) | [Street View](#)

[Suchanzeige schalten](#) | [Günstig umziehen](#) | [Musterbrief Kündigung](#) | [SCHUFA-Auskunft](#)

B

[Hausansicht](#)

[Anbieter kontaktieren](#) | [Merken](#) | [Drucken](#) | [Melden](#) | [Notiz](#) | [Senden](#)

Scout-ID: 14930304 | Objekt-Nr.: 119902

**Kaltmiete:** 247.50 EUR. [Mit lokalem Mietspiegel vergleichen](#)

**Zimmer:** 2.00

**Wohnfläche ca.:** 45.00 m<sup>2</sup>

**SCHUFA-Auskunft:** [Online SCHUFA-Auskunft anfordern](#)

**Hauptkriterien**

Wohnungstyp: Dachgeschoss

Zimmer: 2.00 C

Wohnfläche ca.: 45.00 m<sup>2</sup>

Etage: 3

Schlafzimmer: 1

Badezimmer: 1

Bezugsfrei ab: sofort

Haustiere: Nach Vereinbarung

TV / Internet / Telefon: [Jetzt Verfügbarkeit prüfen](#)

Stromverbrauch: ab 1.500 kWh\* [Stromrechner](#)

\*ImmobilienScout24-Schätzung

Umzugskosten: [Jetzt kostenlose Angebote erhalten](#)

**Kosten**

Kaltmiete: 247.50 EUR

Nebenkosten: + 90.00 EUR

Heizkosten: in Nebenkosten enthalten

**Gesamtmiete:** = 337.50 EUR

Kaution oder Genossenschaftsanteile: € 495,00

Provision für Mieter: Nein

**Bausubstanz**

Objektzustand: Gepflegt

Heizungsart: Zentralheizung

**Ausstattung**

- Laminat
- Bad mit Wanne und Fenster
- WM-Anschluss im Bad
- DSL verfügbar!!! D

**Firma & Kontakt**

**Immobilienmanagement GmbH** ★★★★☆  
33 Bewertungen

Immobilienmanagement GmbH  
[www.immobilienmanagement.de](http://www.immobilienmanagement.de)

[Tel. anzeigen](#) E [Mehr Details](#)

**Nachricht senden** **NEU**  
Persönlich vorstellen

Anrede:  Vorname:  Nachname:

E-Mail:  Telefon:

Gleichzeitig registrieren

AGB und [wichtige Kundeninformationen](#) wurden zur Kenntnis genommen und akzeptiert.

**Besichtigungstermin ist erwünscht.**

[Meine Adresse angeben](#)

[Absenden](#) F

Ihre persönlichen Daten werden absolut vertraulich behandelt und nicht unbefugt an Dritte weitergegeben.

Figure 2.4  
The average broadband penetration of municipalities in 30 quantile bins




Table 2.1  
Descriptive statistics of ImmoScout24 data in the full sample and three estimation samples

	Full Sample					Distance Sample					Above-threshold Sample					OPAL Sample				
	Min	Max	Mean	SD	N	Mean	SD	N	Mean	SD	N	Mean	SD	N	Mean	SD	N			
Monthly rent in €	155	1700	468.29	225.13	1,380,138	531.38	195.48	39,537	508.79	181.06	12,516	361.90	113.70	13,175						
Rent per sqm in €	3	14	6.53	2.06	1,380,138	6.65	1.73	39,537	6.21	1.49	12,516	5.31	0.87	13,175						
Living space in sqm	26	170	71.69	23.40	1,380,138	81.62	25.80	39,537	83.38	26.01	12,516	68.42	18.11	13,175						
Year of construction	1870	2009	1966.55	31.86	1,380,138	1986.07	18.29	39,537	1986.89	18.82	12,516	1973.42	30.94	13,175						
Building age	0	139	41.50	31.88	1,380,138	21.99	18.34	39,537	21.16	18.87	12,516	34.56	30.99	13,175						
Basement	0	1	0.22	0.41	1,380,138	0.26	0.44	39,537	0.24	0.42	12,516	0.13	0.33	13,175						
Elevator	0	1	0.17	0.38	1,380,138	0.07	0.25	39,537	0.04	0.19	12,516	0.06	0.25	13,175						
Garden	0	1	0.16	0.36	1,380,138	0.28	0.45	39,537	0.33	0.47	12,516	0.16	0.37	13,175						
Balcony	0	1	0.67	0.47	1,380,138	0.76	0.43	39,537	0.74	0.44	12,516	0.70	0.46	13,175						
Built-in kitchen	0	1	0.31	0.46	1,380,138	0.45	0.50	39,537	0.42	0.49	12,516	0.16	0.37	13,175						
DSL availability	0	1	0.95	0.07	1,380,138	0.94	0.08	39,537	0.91	0.11	12,516	0.87	0.18	13,175						

*Notes:* The table shows the minimum, maximum, mean, standard deviation and the number of observations of ImmoScout24 properties between 2007-2009 in the full sample and the three estimation samples used in this paper. The Distance sample includes properties in west German municipalities that do not have an own MDF and are connected to the closest MDF from another municipality. The above-threshold sample includes properties in west German municipalities that do not have an own MDF and are more than 4,200 meters away from the MDF they are connected to. The OPAL sample includes properties in east German municipalities that are less than 4,200 meters away from their MDF and do not have more than one MDF.

Table 2.2  
OLS and IV estimates in the distance sample of ImmoScout24 properties

	OLS		Reduced Form		IV	
	(1)	(2)	(3)	(4)	(5)	(6)
DSL availability	0.0057*** (0.0021)	0.0056*** (0.0021)			0.0414** (0.0179)	0.0420** (0.0185)
IV 1			-0.0147** (0.0059)	-0.0144** (0.0059)		
Mobile broadband		0.0004 (0.0006)		0.0002 (0.0007)		-0.0002 (0.0007)
Years	Yes	Yes	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes	Yes	Yes
Regional controls	Yes	Yes	Yes	Yes	Yes	Yes
MDF fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N	37268	37268	37268	37268	37075	37075
R2	0.290	0.290	0.290	0.290	0.277	0.276
clusters	1399	1399	1399	1399	1206	1206
F of excl. instruments					16.15	15.76

*Notes:* Dependent variable is the log of monthly net rents per square meter in EUR. The sample includes properties in west German municipalities that do not have an own MDF and are connected to the closest MDF from another municipality. The instrument is a dummy variable that takes on the value 1 for properties in municipalities that are more than 4,200 meters away from their MDF. Property controls include the age, age<sup>2</sup>, type and condition of the property as well as dummy variables for the availability of a basement, elevator, balcony, built-in kitchen and garden. Regional controls include population, population between 18-65 years, unemployment, share male inhabitants, share of migrants, total area, populated area, street area, and several variables measuring the financial strength of the municipality. Standard errors clustered on the MDF-level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%

Table 2.3  
 OLS and IV estimates in the above-threshold sample of ImmoScout24 properties

	OLS		Reduced Form		IV	
	(1)	(2)	(3)	(4)	(5)	(6)
DSL availability	0.0085*** (0.0020)	0.0069*** (0.0020)			0.0470** (0.0227)	0.0374* (0.0226)
IV 2			-0.0228* (0.0117)	-0.0174 (0.0114)		
Mobile broadband		0.0070*** (0.0014)		0.0071*** (0.0014)		0.0062*** (0.0015)
Years	Yes	Yes	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes	Yes	Yes
Regional controls	Yes	Yes	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N	11829	11829	11829	11829	11826	11826
R2	0.349	0.362	0.348	0.361	0.299	0.331
clusters	184	184	184	184	181	181
F of excl. instruments					19.54	17.84

*Notes:* Dependent variable is the log of monthly net rents per square meter in EUR. The sample includes properties in west German municipalities that do not have an own MDF and are more than 4,200 meters away from the MDF they are connected to. The instrument is a dummy variable that takes on the value 1 for properties in municipalities that cannot be connected to another MDF that is less than 4,200 meters away. Property controls include the age, age<sup>2</sup>, type and condition of the property as well as dummy variables for the availability of a basement, elevator, balcony, built-in kitchen and garden. Regional controls include population, population between 18-65 years, unemployment, share male inhabitants, share of migrants, total area, populated area, street area, and several variables measuring the financial strength of the municipality. Standard errors clustered on the MDF-level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%



Table 2.4  
OLS and IV estimates in the OPAL sample of ImmoScout24 properties

	OLS		Reduced Form		IV	
	(1)	(2)	(3)	(4)	(5)	(6)
DSL availability	-0.0012 (0.0020)	-0.0014 (0.0020)			0.0150* (0.0077)	0.0190** (0.0075)
IV 3			-0.0251** (0.0103)	-0.0329*** (0.0101)		
Mobile broadband		0.0054*** (0.0013)		0.0059*** (0.0014)		0.0051*** (0.0018)
Years	Yes	Yes	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes	Yes	Yes
Regional controls	Yes	Yes	Yes	Yes	Yes	Yes
Counts fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N	12898	12898	12898	12898	12897	12897
R2	0.838	0.838	0.838	0.839	0.835	0.834
clusters	57	57	57	57	56	56
F of excl. instruments					20.32	21.07

*Notes:* Dependent variable is the log of monthly net rents per square meter in EUR. The sample includes properties in east German municipalities that are less than 4,200 meters away from their MDF and do not have more than one MDF. The instrument is a dummy variable that takes on the value 1 for properties in municipalities that are connected to a MDF with OPAL technology. Property controls include the age, age<sup>2</sup>, type and condition of the property as well as dummy variables for the availability of a basement, elevator, balcony, built-in kitchen and garden. Regional controls include population, population between 18-65 years, unemployment, share male inhabitants, share of migrants, total area, populated area, street area, and several variables measuring the financial strength of the municipality. Standard errors clustered on the MDF-level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%

Table 2.5  
Comparison of OLS estimates for each sample in east and west Germany

	Full Sample		Distance Sample		Above-threshold Sample		OPAL Sample	
	West	East	West	East	West	East	West	East
dsl	0.0061*** (0.0022)	0.0005 (0.0015)	0.0057*** (0.0021)	0.0020 (0.0042)	0.0085*** (0.0020)	0.0061** (0.0028)	-0.0028 (0.0041)	-0.0012 (0.0020)
Years	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Regional controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Counts fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	885,735	465,172	37,268	5,610	11,829	1,772	27,878	12,898
R2	0.263	0.245	0.290	0.276	0.349	0.383	0.293	0.838
clusters	3551	1032	1399	350	184	56	1034	57

*Notes:* Dependent variable is the log of monthly net rents per square meter in EUR. The Distance sample includes properties that do not have an own MDF and are connected to the closest MDF from another municipality. The above-threshold sample includes properties that do not have an own MDF and are more than 4,200 meters away from the MDF they are connected to. The OPAL sample includes properties that are less than 4,200 meters away from their MDF and either are connected to an OPAL MDF or do not have an own MDF. Property controls include the log of square meters, age, age<sup>2</sup>, type and condition of the property as well as dummy variables for the availability of a basement, elevator, balcony, built-in kitchen and garden. Regional controls include population, population between 18-65 years, unemployment, share male inhabitants, share of migrants, total area, populated area, street area, and several variables measuring the financial strength of the municipality. Standard errors clustered on the MDF-level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%

Table 2.6  
OLS and IV estimates for west and east Germany in the SOEP sample

	West Germany				East Germany	
	OLS	OLS	IV	IV	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)
DSL	0.0477*** (0.0123)	0.0344** (0.0136)	0.1559 (0.1606)	0.1616 (0.1615)	0.0048 (0.0259)	-0.0153 (0.0283)
Property condition	0.0343*** (0.0122)	0.0218* (0.0120)	0.0396*** (0.0144)	0.0273* (0.0142)	0.0325* (0.0191)	0.0384 (0.0242)
With terrace	0.0738*** (0.0153)	0.0682*** (0.0167)	0.0773*** (0.0156)	0.0734*** (0.0172)	0.0130 (0.0310)	-0.0082 (0.0348)
With central heating	0.1621*** (0.0293)	0.1724*** (0.0343)	0.1552*** (0.0306)	0.1643*** (0.0358)	0.0911 (0.1052)	0.1168 (0.1165)
Property controls	Yes	Yes	Yes	Yes	Yes	Yes
Distance controls	Yes	Yes	Yes	Yes	Yes	Yes
Regional controls	Yes	Yes	Yes	Yes	Yes	Yes
Without city center	No	Yes	No	Yes	No	Yes
N	2149	1808	2149	1808	415	348
R2	0.614	0.626	0.599	0.605	0.615	0.620
clusters	800	785	800	785	138	137
F of excl. instruments			16.41	15.21		

*Notes:* Dependent variable is the log of monthly net rents in EUR. The instrumental variable in Columns 3-4 is the distance between the household and the MDF that the household is connected to. The samples in Columns 2, 4 and 6 are restricted to households that are not located in the city center. Property controls include: living space in square meters, number of rooms, age, recent renovations, recent change of tenants, and the availability of a storage room or a garden. Distance controls include categorical dummies for the distance between the property and the closest: city center, primary school, secondary school, bank, doctor, nursing home, recreation area, kindergarten, restaurant, grocery shops, sports facility, public transport, youth center. Regional controls include: the area, population, average rent, number of firms, number of hospitals, number of residential buildings, share of population with high school diploma, share of school dropouts, share of migrants, and a set of tax related controls on the county level. Standard errors clustered on the municipality level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%

## Appendix

Figure A2.1

The availability of fixed-line broadband access with at least 16 Mbps over the availability of broadband with at least 1 Mbps in relation to the average distance between municipality centroids and MDFs on the state level in 2010

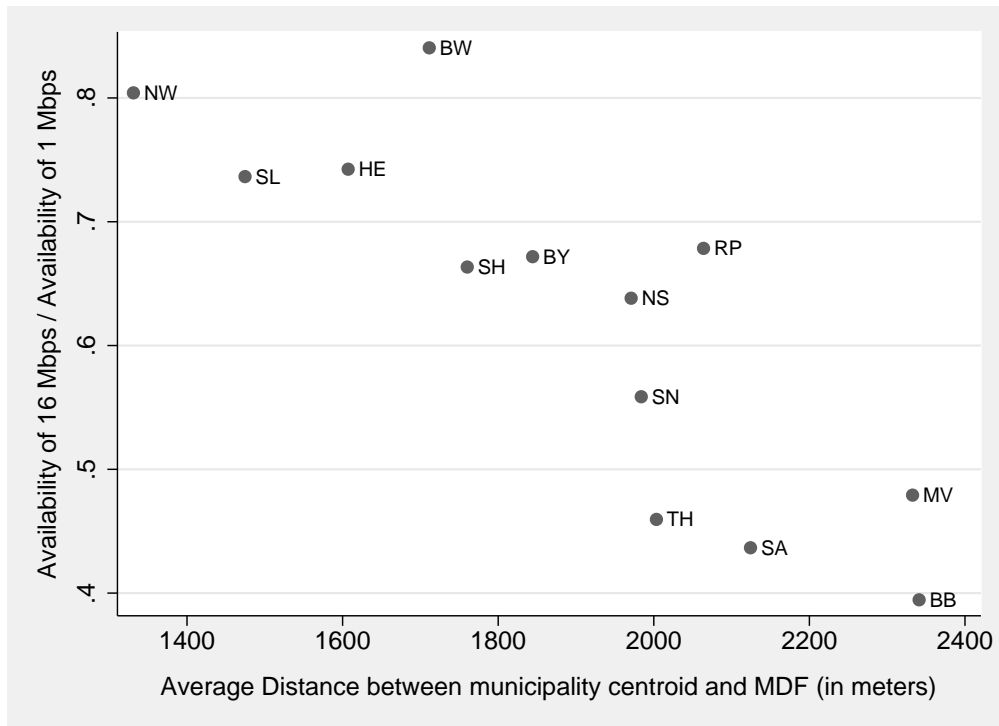


Figure A2.2  
Share of SOEP households with DSL in different distance bins




Figure A2.3  
Fixed-line broadband Internet subscribers in Germany (in million)




Table A2.1  
Descriptive statistics of SOEP 2008 survey data on rented properties

	Min	Max	Full Sample				West Germany			East Germany			Difference
			Mean	SD	N	Mean	SD	Mean	SD	Mean	SD		
DSL	0	1	0.48	0.50	4842	0.52	0.50	0.38	0.49	0.38	0.49	-0.14	
Rent in EUR	101	3559	463.22	222.96	4842	497.47	236.19	391.53	171.52	391.53	171.52	-105.94	
Rent in EUR per squaremeter	1	35	6.34	1.98	4838	6.57	2.17	5.85	1.39	5.85	1.39	-0.72	
Squaremeters	10	310	74.34	27.27	4838	77.64	28.65	67.45	22.64	67.45	22.64	-10.18	
Building age	1	7	3.38	1.56	4591	3.45	1.54	3.22	1.60	3.22	1.60	-0.23	
Number of rooms	1	9	2.98	1.11	4842	3.08	1.18	2.79	0.93	2.79	0.93	-0.28	
General condition	1	4	3.62	0.56	4836	3.61	0.56	3.64	0.56	3.64	0.56	0.03	
Recent renovation	0	1	0.08	0.28	4842	0.09	0.29	0.06	0.24	0.06	0.24	-0.03	
Recently moved in	0	1	0.14	0.34	4842	0.14	0.35	0.12	0.33	0.12	0.33	-0.02	
Has basement	0	1	0.93	0.26	4826	0.92	0.28	0.95	0.22	0.95	0.22	0.03	
Has central heating	0	1	0.97	0.18	4826	0.96	0.19	0.98	0.14	0.98	0.14	0.02	
Has garden	0	1	0.31	0.46	4756	0.37	0.48	0.19	0.40	0.19	0.40	-0.18	
Has terrace	0	1	0.72	0.45	4803	0.75	0.43	0.67	0.47	0.67	0.47	-0.08	
Located in city center	0	1	0.18	0.39	4221	0.16	0.37	0.23	0.42	0.23	0.42	0.06	
Distance to city center	1	6	2.85	1.49	4221	2.83	1.43	2.88	1.62	2.88	1.62	0.06	
Distance to primary school	1	4	1.72	0.88	3939	1.70	0.85	1.75	0.93	1.75	0.93	0.05	
Distance to secondary school	1	4	2.48	1.06	3900	2.53	1.06	2.37	1.04	2.37	1.04	-0.16	
Distance to bank	1	4	1.60	0.82	4229	1.54	0.78	1.73	0.89	1.73	0.89	0.19	
Distance to doctor	1	4	1.88	0.97	4216	1.86	0.96	1.93	0.99	1.93	0.99	0.06	
Distance to nursing home	1	4	2.18	1.06	3882	2.20	1.07	2.12	1.05	2.12	1.05	-0.08	
Distance to recreation area	1	4	1.56	0.86	4184	1.57	0.87	1.53	0.84	1.53	0.84	-0.04	
Distance to kindergarten	1	4	1.63	0.84	3875	1.62	0.83	1.64	0.85	1.64	0.85	0.02	
Distance to restaurant	1	4	1.40	0.68	4213	1.37	0.66	1.45	0.72	1.45	0.72	0.08	
Distance to grocery shops	1	4	1.45	0.75	4228	1.48	0.76	1.40	0.73	1.40	0.73	-0.07	
Distance to sports facility	1	4	1.82	0.88	4133	1.78	0.84	1.90	0.94	1.90	0.94	0.12	
Distance to public transport	1	4	1.14	0.40	4225	1.13	0.40	1.15	0.39	1.15	0.39	0.02	
Distance to youth center	1	4	2.16	1.05	3785	2.14	1.03	2.19	1.09	2.19	1.09	0.04	

*Notes:* Columns 2-6 show the minimum, maximum, mean, standard deviation and number of observations for each variable in a sample of all households in rented non-residential properties without reduced rent. Columns 7-8 show mean and standard deviation for households without DSL and columns 9-10 show mean and standard deviation for household with DSL. Column 11 shows the difference in means between households with and without DSL. Households with properties of less than ten square meters or less than 100 EUR monthly rent are excluded from the sample.

Table A2.2

The association between Internet use intensity and distance to the MDF

	(1)	(2)	(3)	(4)
1,000-2,000 meters	0.0754** (0.0298)	0.0595** (0.0245)	0.0456 (0.0400)	0.0629 (0.0523)
2,000-3,000 meters	0.0269 (0.0421)	-0.0194 (0.0352)	0.0159 (0.0530)	0.0975 (0.0705)
3,000-4,000 meters	-0.0831* (0.0504)	-0.0630 (0.0406)	-0.0628 (0.0577)	0.0119 (0.0765)
4,000-5,000 meters	-0.0141 (0.0594)	-0.0848* (0.0499)	-0.0784 (0.0684)	-0.0256 (0.0996)
5,000-6,000 meters	-0.2243*** (0.0791)	-0.1311* (0.0713)	-0.2168** (0.0916)	-0.1717 (0.1233)
Above 6,000 meters	-0.4012*** (0.1094)	-0.2786*** (0.0950)	-0.2472** (0.1151)	-0.2910* (0.1762)
Individual-level Controls	No	Yes	Yes	Yes
State fixed effects	No	Yes	Yes	Yes
Without agglomeration areas	No	No	Yes	No
Restricted to rural areas	No	No	No	Yes
N	18925	17691	7325	4211

*Notes:* The dependent variable measures Internet use on a scale from 1 (never) to 5 (daily). The independent variables shown are dummies for distance brackets. Distances are measured from the household to the MDF that the household is connected to. The sample in Column (3) is restricted to counties that are outside of agglomeration areas, the sample in Column (4) is restricted to counties in rural areas. Individual-level controls include age, education, gender, migration status, employment, and income. Robust standard errors in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%





## Chapter 3

---

# Online Job Search and Matching Quality

---

An increasing number of people use the Internet to look for new jobs. One reason online job search has become so popular is that it has changed the search process considerably. Employment websites such as Monster.com allow job seekers to access thousands of job offers and use intelligent filter mechanisms to find suitable vacancies. Additionally, online job descriptions provide more detailed information than traditional help-wanted ads in newspapers and magazines. Employers benefit from the better targeting options of Internet job advertisements and are able to screen online applications more efficiently. As a result, the matching process in the labor market has not only become more efficient, but the quality of job matches should be better.

This paper provides first evidence that online job search is associated with higher matching quality. Using micro-level data from the German Socio-Economic Panel (SOEP), I compare employees who found a job online with those who found a job through newspaper advertisements, friends, or other channels. I show that Internet job finders can make better use of their skills, are more content with their work, and believe themselves to have a higher chance of promotion and more job security.

My results indicate that the Internet is an especially valuable job search tool for workers who are distant from the labor market. Job seekers with employment interruptions have significantly better matching outcomes if they find a new job through the Internet. While women with children below the age of 16 generally have inferior results after starting a new job, this negative association is alleviated for those of them who use online job search. I find a similar relationship for job

seekers in rural areas; the disadvantage due to remoteness is remedied if they find a job through the Internet. For workers who were unemployed before they found a new job, I also observe a positive association with online job search, although it is less pronounced than for workers with employment interruptions or for workers from rural areas.

By restricting the sample to workers who found their previous job offline and their current job either offline or online, I can compare the improvement in match quality after a job change conditional on the job search channel. This allows estimations in the spirit of first-differences. Additionally, I apply propensity score matching methods to restrict the comparison group of offline job seekers to those that are most similar to online job seekers. I show that my results hold even if I compare online job seekers only to those who found their job through newspapers and friends. When I compare different search channels to the employment office, I find that only the Internet is associated with significantly higher matching quality. I am able to mitigate numerous selection concerns by robustness tests and providing additional evidence from the German Internet job search market.

The reason the Internet has such a profound impact on the job matching process involves more than the wider selection of job opportunities, better search possibilities, and cheaper access to information. The Internet has introduced new ways of passive job search and allows firms to easily search for applicants. Career-oriented social networks such as LinkedIn and online job boards such as Monster.com allow users to maintain online CVs that can be found by interested employers. Before the advent of the Internet, the direct targeting of talent by firms was feasible only through headhunters and mainly used to fill executive positions. Allowing firms to tap into a large pool of passive job seekers, all of whom provide detailed information about their skills and experiences, results in more informed hiring decisions and contributes to better match quality.

Krueger (2000*b*) was one of the first to note that by reducing the cost of information, the Internet allows workers and employers to learn more about each other and thereby improves the quality of job matches. Autor (2001) points out that due to the Internet, workers and firms are able to consider more potential match partners, which raises the minimum match quality they are willing to accept. The higher match quality in turn leads to higher output and earnings. While Autor acknowledges that better match quality should reduce job separations, he also states that the wider use of on-the-job search has the potential to offset this effect. The increasing popularity of CV databases and career networks such as LinkedIn, which was launched one year after Autor published his article, gives his on-the-job search argument additional weight. Freeman (2002) argues that better job matches

might be the strongest macroeconomic consequence of online job search. Regarding unemployment duration, he suspects that the lower cost of search might ultimately lead to longer search times as workers and firms will consider more possible matching partners. Kuhn (2003) draws on classic partial-equilibrium search models and hypothesizes that by increasing the arrival rate of offers and decreasing search costs, online job search should lead to shorter periods of unemployment and higher-quality job matches.

The empirical literature on Internet job search is mainly concerned with the characteristics of online job seekers as well as with the effect of online job search on unemployment. Kuhn and Skuterud (2004) use U.S. Current Population Survey (CPS) data from 1998 and 2000 to show that once observable characteristics are held constant, Internet job search does not lead to shorter periods of unemployment and might even prolong them. In addition to explaining this finding by stating that the Internet may be an inferior job search tool, the authors raise selection concerns and hypothesize that the longer search time is compensated by improved job quality. Replicating Kuhn and Skuterud (2004) with newer data, Kuhn and Mansour (2011) find that in the period between 2008 and 2009, online job search reduced unemployment duration by about 25 percent. Stevenson (2006) argues that limiting the focus to the unemployed can be misleading as the main effect of online job search could be improved matching outcomes through on-the-job search. Using data similar to those used by Kuhn and Skuterud (2004), Stevenson finds that the Internet has led to higher employer-to-employer worker flows, which could indicate better job match quality for the employed. There is little empirical evidence that online job search reduces net unemployment. Using an instrumental variable approach, Czernich (2012) finds no evidence that broadband Internet affects unemployment rates in Germany. Kroft and Pope (2014) use data from the classified advertisements community Craigslist.com and find that the website's local expansion has to some degree crowded out newspaper advertising but has not had an effect on unemployment rates. But even if there are no effects on unemployment, many studies point out that there should be a substantial effect on matching quality. To my knowledge, however, there has been no study testing this claim directly.

This paper is structured as follows. Section 3.1 describes the data as well as the estimation model. Section 3.2 presents my results and Section 3.3 thoroughly discusses possible selection concerns. Section 3.4 concludes.

### **3.1 Data and Methodology**

There are several different approaches to measuring the quality of a job match. One indirect approach is to use employment duration as an indicator of match quality (e.g., Centeno, 2004). Measures that rely on job tenure assume that "good matches endure" (Bowlus, 1995), which is often but not necessarily true, especially in a relatively rigid labor market like that of Germany. Another approach is to use wages as an indicator of match quality (e.g., Simon and Warner, 1992, van Ours and Vodopivec, 2008). However, the wage of a job changer typically is determined before the employment contract is closed and imperfect information will make it impossible to know the match quality *ex ante*. A way to circumvent this problem would be to consider wage increases in the years after a job change. However, variations in wage are to a large extent driven by supply and demand as well as by other factors that are not necessarily related to the matching quality. Ferreira and Taylor (2011) find that the match quality explains less than 1 percent of wages and Kuhn and Mansour (2011) find no effect of Internet job search on wage growth between jobs.

Therefore, I take a different approach and use subjective matching quality measures from survey data as outcome variables. In the following section I will describe these data as well as the estimation model I use to investigate the relationship between finding a job through the Internet and the matching quality.

#### **3.1.1 Data on Job Search Methods and Matching Quality**

The individual-level survey data I use in this study come from the German Socio-Economic Panel (SOEP), a representative annual panel survey of almost 11,000 households and more than 20,000 individuals. The SOEP was started in 1984 and covers a wide range of topics, including many related to employment. Most importantly for my analysis, people who changed their job are asked not only how they learned about their new job, but also how their new job compares to their former job. One advantage of focusing on job changers instead of also including first-time employees is that workers who had a job before have a reference point with which to compare the new job. Their expectations about how well they can use their skills at work, for example, are likely to be more realistic than those of respondents who have not had much prior work experience. Also, by comparing a new job with an old job both held by the same person, I can perform estimations in the spirit of a first-difference model.

The SOEP asks "How did you find out about your new job?" and provides several answer options from which the respondent must choose one. Respondents choose if they find out about the job in the Internet, through a newspaper advertisement, through friends or relatives, through a private recruitment agency, through a range of different government-run employment offices and job centers, or through none of the aforementioned channels. Respondents can also indicate that they have returned to a former employer.

The variables that measure match quality are constructed based on the question: "How do you view your current position compared to your previous one?" This is followed by a list of sub-questions asking the same question in regard to, for example, "the type of work," "chances of promotion," "work hour regulations," "workload," and "commute," which can be answered by choosing "improved," "about the same," or "worse." There is a separate question that reads: "Are you able to use your professional skills and abilities today more, about the same, or less than in your previous position?" The variables are coded 1 when the answer is "improved" and 0 otherwise. Alternatively, I estimate an ordered logit model with all three answer choices and find results similar to those from the binary choice model.<sup>1</sup>

The relatively small number of observations makes it necessary to pool all observations between 2000 and 2007. Unfortunately, later years cannot be used in the analysis because the SOEP waves after 2007 do not include the relevant questions concerning the assessment of the new job compared to the previous one. By limiting the focus to job changers between 2000 and 2007, my sample is reduced to about 2,000 observations per year. As shown in Table 3.1, the share of people who found a new job through the Internet increases from less than 1 percent in 2000 to more than 6 percent in 2007. These relatively low numbers are due to the fact that I do not include in the sample the young and often Internet-savvy workers who found their first job through the Internet. The numbers also do not reflect how many job seekers actually used the Internet at any point during their job search, but only more conservatively reflect the number who learned about the job through the Internet and chose to sign a contract with the employer. Thus, it is likely that the estimates represent a lower bound of a potentially stronger association.

Table 3.2 shows the sample means by the channel used to find a new job. Online job finders in the sample are on average 32.5 years old, which is slightly younger than workers who found a new job through other channels, even though the difference is not significant. There are also more men among the online job finders compared to those who used the newspaper, for example. A large

---

<sup>1</sup> Results are available upon request.

share of employees who returned to their former employer are women, a finding probably driven by mothers who were on parental leave. Internet job finders are on average slightly better educated than those who found a job in the newspaper. The share of formerly unemployed job changers who used the Internet is slightly higher than the one of employees who found a job through friends or newspapers. Although some of these differences between groups are interesting, few of them are very large or even statistically significant.

### 3.1.2 Estimation Strategy

The advantage of the SOEP data is that I can focus on job changers and compare how the same person evaluates his new job compared to his previous job. This allows estimations in the spirit of first differences. By comparing two jobs of the same person, I implicitly control for job searcher characteristics that could bias results of a conventional cross-sectional regression of current matching quality measures on job search methods.

The estimation model has the following form:

$$M_{itsc} = \alpha + \beta internet_{itsc} + \mathbf{X}'_{itsc} \gamma + year_t + industry_s + county_c + \varepsilon_{itsc} \quad (3.1)$$

where  $M_{itsc}$  are the subjective matching outcome variables of a job changer  $i$  in year  $t$ , industry  $s$ , and county  $c$ . Specifically, they are an employee's evaluation of: ability to apply own skills, satisfaction with the type of work, career perspectives, job security, social benefits, workload, commute, and working hours. The dependent variables always indicate how a person evaluates the new job compared to the prior job and can therefore be interpreted as the change in evaluations between two jobs. *internet* is a dummy indicating whether a person found the job through the Internet.  $X$  is a vector of individual-level covariates, including gender, age, migration status, education, number of job changes between 2000 and 2007, and a dummy indicating whether the person was unemployed during the last year. To limit the risk of merely observing a correlation based on Internet usage in general, I include a dummy that indicates the availability of the Internet in the household. Additionally, I include year, industry, and county fixed effects. Since the dependent variables are binary, I estimate a probit model with robust standard errors.

As shown in Table 3.1, the number of job seekers who found their new job in the Internet is relatively low compared to the total number of job changers. One could therefore be worried that at least part of the job changers who used other channels than the Internet are hardly comparable

to the Internet job seekers. In order to restrict the comparison group, I apply different propensity score matching methods. Given the observed characteristics of a job seeker, the propensity score gives the conditional probability of finding a job through the Internet (see Rosenbaum and Rubin, 1983). One advantage of the propensity score is that its estimation does not depend on linearity assumptions for the dependent variable (Heckman, Ichimura and Todd, 1998). With the help of the propensity score, I can evaluate which non-Internet job seekers are most comparable to the Internet job seekers in the sample. By restricting my comparison group to a subset of job seekers that are similar to Internet job seekers, I am able to alleviate bias due to systematic observable differences between the two groups.

There are several methods to match observations based on their propensity score.<sup>2</sup> One of the most straightforward ways to find a match for every Internet job seeker is to choose the observation with the smallest propensity-score distance in the group of non-Internet job seekers. This nearest-neighbor method leads to a comparison group that is very similar to the group of interest, but comes at the cost of precision because of potentially low observation numbers. In order to increase the sample size, I also apply an algorithm that selects the 5 nearest neighbors for every Internet job seeker. The 5-nearest neighbor matching reduces the variance by using more information, but results in potentially larger bias due to on average poorer matches. To use even more information on the comparison group, I finally apply Epanechnikov kernel matching which uses the weighted average of all individuals from the comparison group to construct the counterfactual.

Figure 3.1 shows the distribution of propensity scores for Internet job seekers and for the comparison group of job changers that found their job through other channels. The light grey bars show observations that are excluded from the sample when 5-nearest neighbor matching is applied. Note that only 3 observations are excluded from the sample because they fall out of the common support region. 5,676 observations from the comparison group are excluded due to their propensity score distance to Internet job seekers. Among the Internet job seekers, 3 observations from the top of the propensity score distribution are excluded from the sample.

---

<sup>2</sup> see Smith2001 or Dehejia and Wahba (2002) for a comparison of different propensity score matching algorithms

## 3.2 The Association of Internet Job Search and Matching Quality

In this section, I begin by presenting results from the baseline model of Equation 3.1 and then turn to results from propensity score matching. I further explore effect heterogeneities and find that Internet job search seems to be especially valuable for job seekers that are distant to the labor market. Finally, I compare job matching outcomes of the Internet with those of other job search channels.

### 3.2.1 Baseline Results

Table 3.3 presents the results of regressing different job matching outcome variables on the Internet search dummy as well as other covariates. All reported coefficients are probit marginal effects. In the first column, the positive and significant Internet coefficient indicates that online job seekers are more than 6 percent more likely to report that they use their skills better in their new job. They are also significantly more likely to be satisfied with the type of work they do, as the high Internet coefficient in the second column shows. The dependent variable with the highest Internet coefficient is the perspective variable in the third column. It shows that online job seekers are more than 8 percent more likely to have a better chance of promotion in the new job. Finding a job online is also associated with better job security, as shown in Column 4. Surprisingly, in Column 5 we see that also social benefits,<sup>3</sup> such as holiday pay, company pension schemes, or corporate child care, are significantly better for online job seekers.

In the last three columns of Table 3.3 we see results for dependent variables that are not significantly correlated with online job search. Column 6, indicating satisfaction with workload in the new job, has an Internet coefficient that is positive but below 1 percent and insignificant. In Column 7, Internet even has a very small negative coefficient on the satisfaction with the commute. This could mean that online job seekers are more likely to find a job that is farther away from home than the previous job. A possible interpretation of this finding is that the Internet opens up job opportunities outside regional boundaries and thereby increases work mobility. Online job search is also barely associated with work time, as shown in the last column. Unlike the other dependent variables in the Table men are significantly less likely to improve either their commute or working

---

<sup>3</sup> *betriebliche Sozialleistungen*



time compared to women. This might indicate that these dimensions are less important to men when they change jobs.

In most cases, the coefficients of the control variables have the same sign across dependent variables even though their size differs. Being male is positively associated with the first five outcome variables. This means that, on average, male job seekers evaluate their new job better than female job seekers do. One explanation could be that men either obtain better jobs or are more optimistic about a recent job change. The opposite can be observed for older job changers compared to younger ones. A possible explanation could be that older job changers were already better matched in the previous job and have therefore less room for improvement after a job change. Since our outcome variable only measures improvement compared to the previous position, this could lead to a negative association between the outcome variables and age. A similar negative correlation can be observed for people with a migration background as well as for job changers who experienced a period of unemployment before finding a new job. The more highly educated a person, the better the matching outcome.

Note that the Internet coefficient is positive for all dependent variables except for commute. The reason we see quite small and insignificant coefficients on the workload and working time variables might be that these are relatively poor measures of matching quality. When a person switches to a new job that is different from his or her previous one, the workload might initially be greater than at the old job as the person needs to become familiar with new tasks and processes. It is therefore possible that a higher workload could signal a good match in some cases but a poor match in others. The four variables that are much more clearly measures of matching quality-skill use, work type, perspective, and job security-all have high and significant Internet coefficients. In the remainder of this paper I chiefly discuss the first four dependent variables because I believe they are most relevant for assessing matching quality. In the next section, I analyze heterogeneity effects on the skill use variable, which is arguably the most interesting measure of job match quality.

### **3.2.2 Results in Matched Samples**

Table 3.4 presents estimates from regressions on samples that are matched with three different matching algorithms. Every coefficient is from a separate regression that uses the respective matching weights and includes all control variables from the baseline results in Table 3.3. The same covariates were used for the estimation of the propensity score. The upper panel shows results of Epanechnikov kernel matching that reduces the number of observations to 6,430 for the first outcome

variable. Although about 35 percent of the observations in the comparison group are excluded from the sample, the coefficients of finding a job online remain robust for all outcome variables. For job security, the coefficient becomes slightly smaller but remains statistically significant at the 5 percent level.

The middle panel shows regressions in a 5-nearest neighbor sample that reduces the number of observations to about 1,700. For the skill use, work time, and perspective outcome variables, the coefficients are only marginally different from the coefficients in the full sample. The coefficient for job security decreases from 4.55 percent in the full sample to 2.85 percent in the 5-nearest neighbor sample. The lower panel presents results from nearest neighbor matching. Despite a strong reduction in the number of observations, the coefficients of the first three outcome variables remain statistically significant. For skill use, the coefficient even increases significantly to more than 10 percent. For job security, the coefficient decreases further to about 1.5 percent.

The results of Table 3.4 show that the large difference in the number of Internet job seekers and job seekers that use other channels does not seem to bias the estimates of finding a job online on the skill use, work type and perspective outcome variables. Even after selecting only the observations that are most similar to Internet job seekers, the estimates remain robust. After restricting the sample to the closest matches, the positive association between finding a job online and better job security is reduced to a third and also loses statistical significance.

### **3.2.3 Effect Heterogeneity**

To this point, we have been concerned with the average association of online job search and matching outcomes among all job changers. Table 3.5 shows that the strength of this association varies depending on the subgroup to which the job changer belongs. Each line in Table 3.5 represents one least squares regression with "skill use" as the dependent variable and the same control variables as in Table 3.3. Additionally, each regression contains the variable in the lead column and an interaction effect of this variable with the "found via Internet" dummy. The first column shows the association between finding a job online and being able to use personal skills better in the new job for workers who do not belong to the group described in the lead column. Column 2 shows the main effect of the variable in the lead column on the ability to use own skills. The third column reports the estimates for the interaction term of Internet job search and the respective variable in the lead column.

Workers who just reentered the employment market are 10 percent less likely to feel that they can use their skills better in their new occupation, as shown in the second column of the first row. These workers were not unemployed before they found a new job. Although the exact reason for their employment interruption is unclear, the high proportion of women in this group points in the direction of parental leave. Other possible reasons for such career breaks include educational leaves, national service, volunteer work, travel, or rest. The literature on employment interruptions argues that a worker's human capital stagnates or even decreases during career breaks, with the exception of educational leaves<sup>4</sup>. The skills acquired in school and during previous occupations become increasingly outdated and depreciate during employment interruptions. According to Williams (2000), even career breaks due to self-employment can have adverse effects as sector-specific human capital decreases over time. Mincer and Polachek (1974), who underline the importance of work history in human capital models, point out that during periods of childbearing, erosion of market skills might lead women to revise their expectations of and commitment to employment. The strong negative coefficient for women with children in the second row of Table 3.5 could be interpreted as support for this idea. Besides the human capital effect, there is a signaling effect induced by career breaks. Employers could interpret an employment interruption as a sign of low commitment or reliability. Consequently, they might be reluctant to offer jobs involving much responsibility to workers with career breaks. Along with skill depreciation, this could explain the negative association of reentry and being a mother with the outcome variable.

The third column of Table 3.5 shows a significantly positive interaction effect for workers who just reentered the employment market and found their job online. The same reversal takes place for women with children. This could mean that the Internet is an especially valuable job search tool for workers with employment interruptions. One explanation for this finding could be that the negative signaling induced by career breaks is less severe when the job is intermediated through the Internet. Another, probably more convincing, explanation is that the Internet is especially important for workers who are more distant from the labor market. For example, women who are caring for their children instead of engaging in formal employment are less likely to hear of current employment opportunities in the organization or industry in which they previously worked.

The third row seems to confirm the hypothesis that distance from the market matters. For those job changers who live in a county that has a population density below average, denoted as a rural county in Table 3.5, the Internet interaction term is sizable and significant. This finding relates to the "death of distance" hypothesis Cairncross (1997). Cairncross argues that modern

---

<sup>4</sup> For a recent overview of the literature on career breaks see Theunissen et al. (2011).

telecommunication networks will improve rural areas' access to larger markets. The disadvantage of job seekers in remote areas is alleviated by the Internet, which opens up new supraregional employment opportunities. Note that in all specifications of Table 3.5, I control for Internet availability in the household. The positive coefficient for online job seekers in rural areas is therefore more than a sign of being better connected due to Internet access; it indicates that online job search makes a difference for those who are distant from urban centers.

Workers who were unemployed before they changed jobs do not seem to benefit more from online job search than other job changers. There are several reasons why job changers who re-enter into employment for other reasons benefit more from online job search than the unemployed. Stevenson (2006) argues that the Internet leads to an increase in on-the-job search, which reduces transitions from employment to unemployment as workers can more easily find a new job online before their current job terminates. This hypothesis seems plausible in light of the passive job search opportunities enabled by the Internet. However, it implies that those who become unemployed are negatively selected with respect to their ability to use online job search to their advantage. Those that exit the labor market temporarily for other reasons do not have such selection bias. In other words, someone who becomes unemployed nowadays might not have the capabilities to benefit from online job search in the first place.

This incapacity could be explained by a lack of exposure to the Internet at the former workplace. Krueger (2000a) argues that the digital divide with regard to race might be partially caused by a underrepresentation of minorities in positions that use computers. Similarly, unemployed job seekers might be less successful with online job search because they were less likely to use the Internet at their former workplace. Due to this lack of expertise, they might use inferior Internet search tools. The problem for the unemployed could therefore be a lack of complementary skills that are necessary to use the Internet to their best advantage.

It is interesting that job seekers below the age of 30 are not benefitting disproportionately from online job search. This indicates that there is no digital divide based on age when it comes to using online job search tools. Although, on average, young people are able to use their skills better after a job change, this is not due to the method they used to find that job. Another concern often raised in context of the digital divide debate is that minorities are disadvantaged when it comes to use of the Internet (e.g., Hoffman and Novak (1998), Fairlie (2004)). While we do observe a negative association between having a migration background and the outcome variable, online job search is not less effective for migrants. In fact, the respective coefficient in Table 3.5 is relatively large and

positive, but not significant at the 10 percent level. Similarly, workers with tertiary education do not benefit more than those with an average level of education from online job search. This result can be interpreted as implying that the higher educated are not necessarily the main beneficiaries of online job search.

### **3.2.4 How Internet Job Search Compares to Other Search Methods**

The previous sections compared the Internet with all other means of finding a new job. But what if the positive correlations shown in Table 3.3 are mainly driven by the comparison with job search tools that lead to especially poor matching results? One channel that could lead to mediocre matching results is public employment services. For example, Holzer (1988) finds that searching for a job through family, friends, and newspapers is associated with a higher probability of receiving an offer than searching through the state employment agency. Clark (1988) shows that the retention rates on jobs facilitated by the public U.S. Employment Service (USES) are lower than those facilitated by other intermediaries. Using data from Portugal, Addison and Portugal (2001) also find that the public employment service is associated with shorter job retention. Additionally, they show that rewards for observable characteristics of job seekers are smaller in jobs found through the public employment service. One possible explanation for these differences could be that public employment services have less incentive to find good matches than do private intermediaries, as argued by Zweifel and Zaborowski (1996).

Therefore, I analyze how different job search channels compare to finding a job through federal and local employment offices as well as through so-called Personalserviceagenturen (PSAs), which are temporary employment agencies attached to employment offices. Table 3.6 shows that the Internet is the only channel with significant positive coefficients across all outcome variables in comparison to the employment office. Although most job changers find their job through friends, acquaintances or family, this channel is not positively associated with better job matching outcomes. One possible reason could be that if someone finds a new job through personal connections, the formal job screening process, which would normally assure a good or at least reasonable match, is not taking effect. The significantly positive coefficient for job security does not contradict this hypothesis. In fact, someone who finds a job through a friend who works for the same organization might feel that the job is especially secure since it is protected by the friend. Another explanation is offered by Loury (2006) who argues that job seekers turn to informal search channels like family

and friends as a last resort and have few alternative choices. Private job agencies are also unable to outperform the employment office, as shown in the fourth row. The "other" category has positive coefficients for most of the matching outcomes. One explanation could be that headhunters and other personal matchmakers fall in this category. The newspaper coefficients are also positive, although insignificant and much smaller than the Internet coefficients. One could argue that people who find a job through the employment office, a private job agency, or by some other undefined means are not of primary interest for the analysis. After all, these channels are very different from job search on the Internet and job seekers who use private job agencies, for example, might have different characteristics from those who use the Internet. To test this argument, I exclude all job changers who used channels other than the Internet, friends, and newspapers and repeat the estimations from Table 3.3. In Table 3.7, we see that the Internet coefficients are considerably lower for all outcome variable. As in the nearest neighbor matching sample, the association between finding a job through the Internet and better job security is considerably weaker and loses statistical significance. For the other outcome variables, the coefficients remain significant at the 5 percent level, at least. This robustness of the Internet coefficients demonstrates that online job seekers are better matched not only compared to all other search methods taken together but also compared specifically to the most similar channels, namely, newspapers and friends.

### **3.3 Discussion of Potential Selection Issues**

The last section demonstrated the robustness of the results to a wide range of controls and sample restrictions. However, there are some obvious selection issues that could bias the results. Possible problems could arise if Internet job seekers are fundamentally different from workers they are compared with. A second cause of concern is that online job seekers search differently from job seekers who use other channels, and are therefore able to find better matches. A third issue is that the kind of companies that use online job tools could be different from companies that advertise in newspapers or through other more traditional channels. The following section tackles these concerns.

#### **3.3.1 Systematic Differences in Job Seeker Characteristics**

Above, we saw that the associations between finding a job online and better job matching outcomes are not driven by observable characteristics such as age, job position, or industry. But this does not

necessarily mean that there are no systematic differences in job seeker characteristics that could potentially bias our results. There are plenty of reasons why one could assume that workers who use the Internet for job search are different from workers who prefer to read job advertisements in newspapers. For example, online job seekers could be generally more open to new technologies and adapt better to technological changes. This could, in turn, be a characteristic valued by employers and correlated with better matching outcomes.

The problem of this selection issue is less severe than it might appear at first because I basically compare two jobs held by the same person. In a setting that is comparable to a first-difference estimation, I focus on workers who have changed jobs and compare their current job to their former job. It might well be the case that workers who find a job on the Internet are better able to take advantage and adapt to change and thus will always have a more positive perspective on their careers than offline job seekers. But as long as this personal characteristic is constant over time, there is no reason to believe that the same person would have more optimistic perspectives in one job compared to another job for reasons that are unrelated to the job itself.

Nevertheless, there are certain selection criteria that have the potential to bias my results. If, in general, online job seekers view a job change more positively than offline job seekers, they could also believe that their new job suits them better simply because they changed it recently. Although testing for personal attitudes toward a new job is difficult, there are some questions in the SOEP that can be used as indicative evidence for a possible systematic difference in attitudes between online and offline job seekers. Unfortunately, most of these characteristics are only asked periodically in SOEP waves. For this reason, I can only use a subsample of job changers and cannot exploit the variation over time. In one question of the SOEP, respondents are asked whether the statement "When I think about the future, I'm actually quite optimistic" applies to them. The first column of Table 3.8 shows that job changers who found a new job on the Internet might be slightly more optimistic about the future, but not to a statistically significant extent. Given that the scale on which the dependent variables in Table 3.8 are measured goes from 0-10, the coefficient is very small.

A second question is concerned with overall life satisfaction and asks: "How satisfied are you with your life" (scale 0-10). Column 2 of Table 3.8 shows that if anything, Internet job seekers are slightly less satisfied with their life than others. This is remarkable and underlines that the job matching outcomes used in this analysis capture something that is much more specific than general life satisfaction.

The SOEP also incorporates rough measures of the so-called Big Five personality traits: openness, conscientiousness, extraversion, agreeableness, and neuroticism. These personality traits are fairly stable over time and found to change very little after the age of 30 (Terracciano, McCrae and Costa, 2010). For our purposes, openness seems to be the most interesting of these traits. In the SOEP, openness is measured in terms of being original and coming up with new ideas as well as having an active imagination<sup>5</sup>. The third and fourth columns of Table 3.8 show that there is no statistically significant difference between online and offline job seekers along these dimensions.

Although there are good reasons to believe that selection on unobservables is not a primary concern in my analysis, I can use an additional test to eliminate possible selection bias. If people who find a job online are systematically different from others, this difference should not only affect the variables where I see a significant positive association, namely, usage of skills, satisfaction with type of work, career perspectives, and job security; the difference between online and offline job seekers should also affect other variables, such as satisfaction with the working time. Let us assume, for example, that online job seekers view their new job more positively than others because of differences in personal attitudes. Then the generally better assessment of the job should make these workers more content with their tasks but also with their working time. Since there are few objective reasons why online job seekers should have more convenient working times, I can use the assessment of working time as a reference point for all of that person's judgments. By including the "working time" variable in the specifications of Table 3.9, I control for a possible selection effect on the evaluation of a given person. While the association of working time with the outcome variables is highly significant, the coefficients of the Internet variable remain stable.

### 3.3.2 Selection into Search Intensities

Another selection concern is that online job seekers may be spending more time looking for a job than do offline job seekers. If workers who use the Internet for job search generally search more intensively, this could to some part explain why they are better matched if they decide to change jobs.

To test whether people who use the Internet for job search are more actively searching than others, I use additional information from the SOEP. Respondents who changed their job are asked:

---

<sup>5</sup> The original questions from the SOEP read: "I see myself as someone who is original, comes up with new ideas" and "I see myself as someone who has an active imagination".



"Were you actively looking for a job when you received your current position, or did it just come up?" Table 3.10 shows that the active search coefficient is positive and highly significant, indicating that people who search actively have considerably better matching outcomes than people who find a job by chance. The coefficient of online job search, however, remains positive and significant for the first three of the four dependent variables in Table 3.10. Only for the job security outcome the Internet coefficient becomes insignificant, similar to what we saw in some of the regressions with restricted comparison groups. The robustness of the other three coefficients indicates that the difference in matching quality is not only driven by searching actively or not.

It is also useful to think about which job changers are likely to have searched more intensively than others. Finding a job through family and friends, for example, might not indicate a very high search intensity. Indeed, Table 3.6 shows that finding a job through family and friends is not generally associated with superior matching quality. For other channels, like newspapers or job agencies, it is less clear that the job search is less intense than on the Internet. Nevertheless, Table 3.6 shows that the Internet is the only channel that is consistently associated with significantly better matching quality compared to the employment office.

The Internet might often be used for job search in combination with other channels. For example, someone might read a newspaper advertisement for a job and then browse the website of the potential employer to learn more about the company and the position in question. In this case, the job seeker benefits from some of the advantages of online job search, namely more information and transparency. However, the relevant survey question in the SOEP reads: "How did you find out about your new job?"<sup>6</sup> The person in our example would probably indicate that he found out about the job through the newspaper. Thus, I might underestimate the association between online job search and matching quality because some of the supposed offline job seekers also used the Internet at some point during the job search. It is less likely that the argument also works the other way around and Internet job seekers look for information about a vacancy in newspapers after finding out about it online.

One reason why it appears unlikely that online job seekers are more serious about their job search is the low search cost on the Internet. Looking for a job in a newspaper, for example, is much more costly than using online job boards. First, newspapers themselves cost money. Second, it is more difficult and more time consuming to find advertisements in a newspaper that match own qualifications. And third, compiling a classic job application, including postage, is more expensive

---

<sup>6</sup> "Wie haben Sie von dieser Stelle erfahren?"

than filling out an online application form or sending an email. These costs are one reason why the number of postal applications has declined steadily over the last couple of years while the use of electronic applications has increased over time, as shown in Figure 3.2.

In addition to being a less expensive search method, the Internet also offers passive job search opportunities, as discussed in the introduction. Today, over 175 million people maintain online CVs on the largest professional network LinkedIn,<sup>7</sup> which was launched in 2003 and already had more than 15 million members in 2007. Xing, a German competitor of LinkedIn, had almost 5 million members in 2007.<sup>8</sup> Online CVs allow recruiters and headhunters to search for job candidates. Also online job boards such as Monster offer the opportunity to upload CVs that can be viewed by recruiters. The whole idea of passive job search clearly contradicts the argument that online job seekers generally spend more time or look more intensively for new job opportunities. As Figure 3.3 illustrates, a majority of people who are interested in career opportunities already use online career networks and CV databases. This shows that online job search can be almost completely effortless, which is not true for most other job search channels.

### 3.3.3 Selection of Advertised Jobs

In the previous sections, I have not yet addressed the employer side. My results could be biased if only a certain kind of company uses the Internet for recruitment purposes. I would overestimate the association between online job search and matching outcomes if the jobs advertised only were better than other jobs. For example, companies that advertise vacancies through the employment office might be less attractive than companies who use online job boards. However, we have seen in Table 3.7 that the results remain stable if I exclude employment offices and job agencies from the analysis. There is also evidence that many high-end jobs are not advertised over the Internet but only through other channels. Specialized headhunters and HR consultancies definitely play an important role in filling executive positions. Many people also believe that high-salary jobs are more often advertised in newspapers. Some companies prefer newspaper advertisements because they look more expensive and signal the value that the company puts on the position. This kind of selection, however, would lead to an underestimation of the association we observe.

Figure 3.4 provides more evidence that companies that advertise online are not necessarily the more attractive employers. It is striking that 40 percent of the companies who advertise on Monster

---

<sup>7</sup> as of August 2, 2012 (see <http://press.linkedin.com/About-Us>)

<sup>8</sup> <http://www.statista.com/statistics/263570/member-numbers-of-social-network-xing/>

are temporary work companies. Usually, this type of company is not a job seeker's most favorite and often pays less than other employers. At Stellenanzeigen.de, another popular German job board, the share of temporary work companies is significantly lower but still twice as high as the share of companies listed on the German DAX stock index. The dominance of the generally less attractive temporal work companies in online job boards would again lead to an underestimation of my results.

Companies choose the advertising channel that they expect will be most effective at attracting appropriately qualified applicants. Different types of jobs are therefore advertised through different channels, a fact that gives rise to another selection concern: How do job positions advertised through the Internet compare to job positions advertised offline? Figure 3.5 shows how well the search results at Monster.de match job titles used as search terms. Almost 80 percent of the search results exactly match the request. There seems to be a tendency that jobs requiring high qualifications, like general manager, HR director, lawyer, or engineer, obtain inferior results compared to more mid-range jobs like controller or project manager. Although, both job advertisements and job searches are endogenous, Figure 3.5 indicates that online job boards are less well suited for top jobs. If the Internet serves as a better channel for mid-level jobs and top jobs are more often advertised through other channels, my results would be underestimated.

There is anecdotal evidence that although online advertisements attract a great quantity of applications, many of them are of lower quality. This could be related to the low costs of application discussed above. From an employer perspective, this problem becomes more severe, the higher the desired qualifications. Figure 3.6 shows the results of a study in which identical job advertisements for the position of procurement director were placed in eight German newspapers and eight German online job boards. The applicants were subsequently rated according to their qualifications. Overall, the number of applications in response to the online advertisements was more than 2.5 times the number of applications in response to the print advertisements. However, more than 50 percent of the Internet applicants were not all qualified for the job, whereas this share was considerably lower for newspaper applicants. Although these numbers are purely descriptive and by no means representative, they do add some weight to the view that online recruitment is a bit more about the quantity and less about the quality of applicants. The numbers also indicate that the better job matches achieved through online job search come at a cost. In the language of network theory, online job search facilitates network formation, but increases costs associated with network clearing. As a result, we might see fewer high-qualification jobs advertised online, which again would lead to an underestimation of the results.

### 3.4 Conclusion

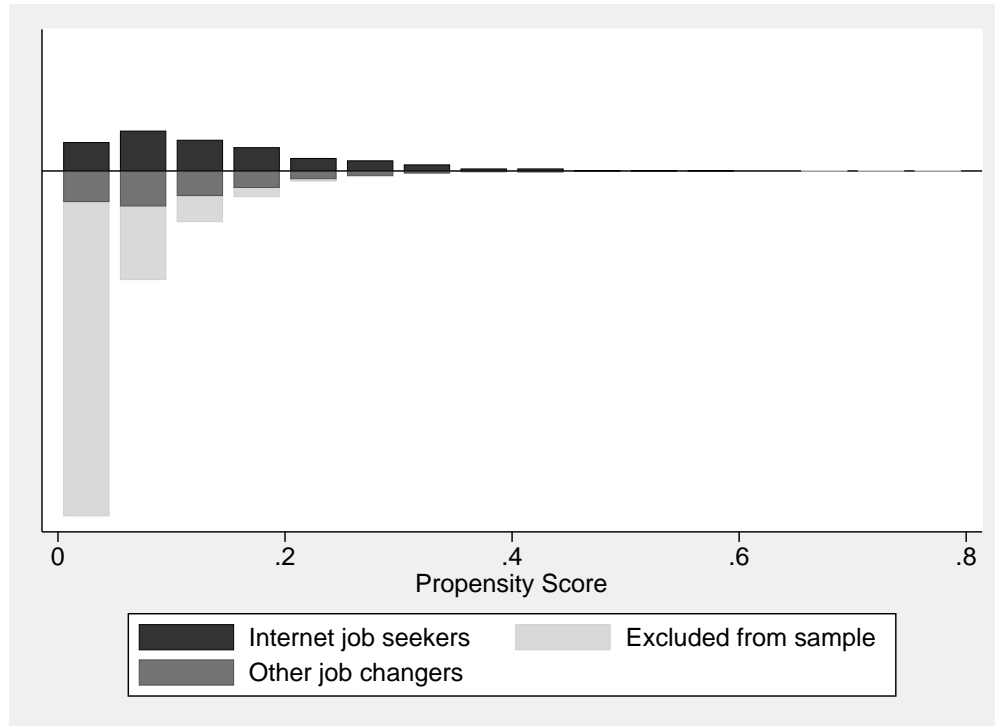
This paper investigates whether online job search is associated with better quality job matching. The question of how the Internet affects unemployment duration and other labor market outcomes is much studied in literature, but this paper provides the first empirical evidence as to the quality of resulting job matches. I find that matching outcomes of online job seekers are superior along several dimensions, including making better use of own skills, being more content with the type of work, having higher chances of promotion, and enjoying greater job security. These results are not driven by comparing the Internet with inferior search channels like the employment office. Online job search is associated with better matching quality even if it is directly compared to searching newspapers or asking friends and family for help. My results avoid bias from many possible sources of selection. As I focus my analysis on workers who found their previous job offline and their current job either offline or online, I can compare two matching outcomes for the same worker by using retrospective data. Additionally, I tackle several selection issues with robustness tests and provide some descriptive evidence to alleviate remaining selection concerns. Even though I am able to rule out the most obvious threats to a causal interpretation of the associations presented in this paper, more work is needed to identify a clear causal relationship between online job search and matching quality.

The results indicate that the Internet is an especially important tool for job seekers distant from the labor market. Workers with employment interruptions are particularly likely to be well matched if they used the Internet to find their new job. Online job search also seems to play an important role for mothers with children. As gender inequality remains an issue in many labor markets, it is important to know that the Internet can alleviate possible negative consequences of a maternity leave. The results also show that job seekers in areas with lower population densities are better matched when they find their job online. This finding has policy implications with regard to the expansion of broadband Internet in rural areas. It is remarkable that online job search seems to compensate for many of the disadvantages suffered by job seekers who are distant from the labor market. Formerly unemployed job seekers also benefit from online job search, but not to the same extent as job seekers who reenter the labor market or who come from rural areas. One possible reason for this finding is that the unemployed lack the necessary complementary skills to use the Internet to their advantage. If this is indeed the case, it might be worthwhile to train unemployed job seekers in new technologies during their job search process. Further research is needed to better understand the relationship between online job search and Internet-related skills.

## Figures and Tables

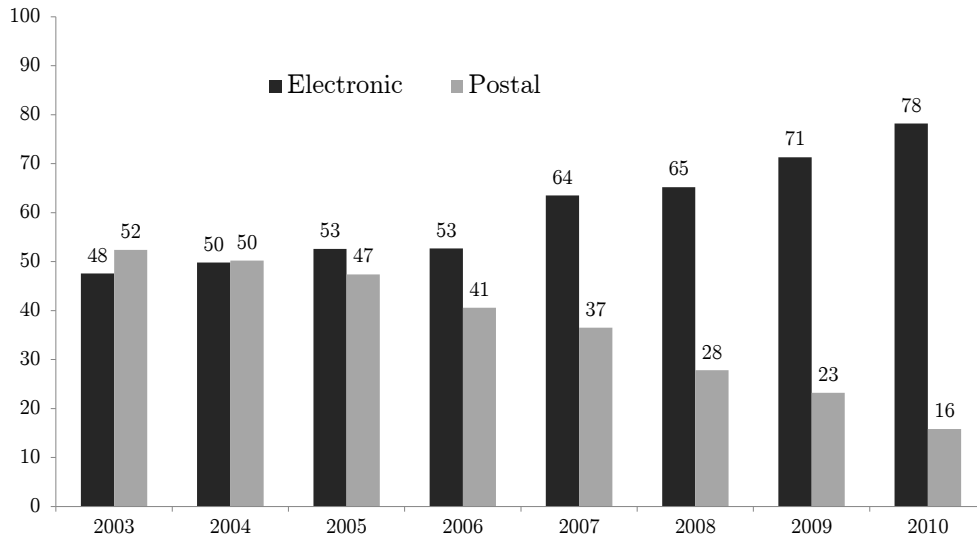
Figure 3.1

Distribution of propensity scores for individuals who found their job online and individuals who found their job through other channels



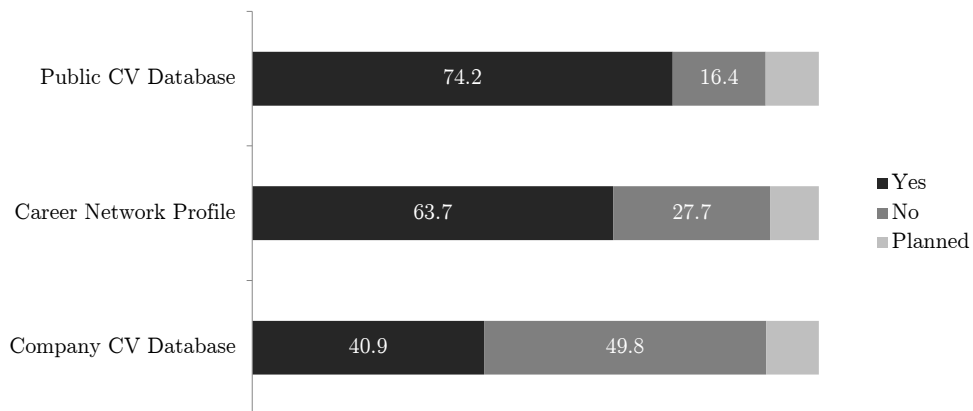
Notes: Results from five-nearest neighbors matching with replacement. Observations that are excluded from the sample include 3 observations that are without common support, 5,676 observations in the comparison group that are not matched to Internet job seekers, and 3 observations in the group of Internet job seekers that are not matched with observations from the comparison group.

**Figure 3.2**  
 Preference of job searchers for electronic and postal applications in Germany over time



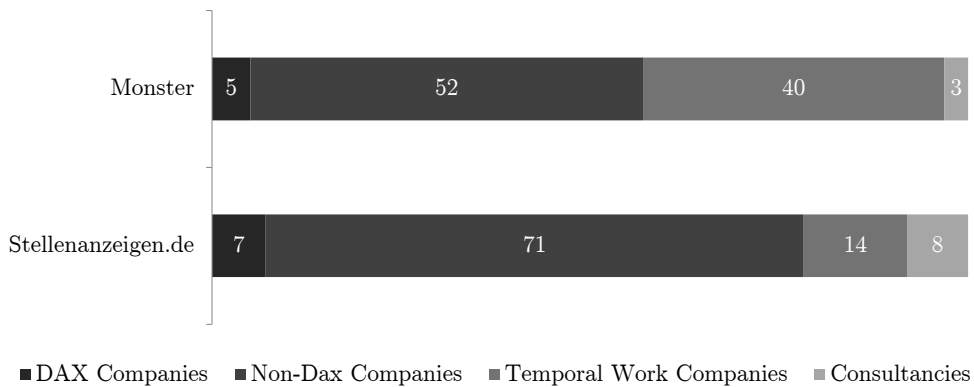
Source: Bewerbungspraxis 2011, Centre of Human Resources Information Systems (CHRIS). Based on 10,227 individuals interested in career opportunities.

**Figure 3.3**  
 Usage of CV databases and online career networks for passive job search in Germany




Source: Bewerbungspraxis 2011, Centre of Human Resources Information Systems (CHRIS). Based on 10,227 individuals interested in career opportunities.

Figure 3.4  
Company types that advertise on two major German online job boards



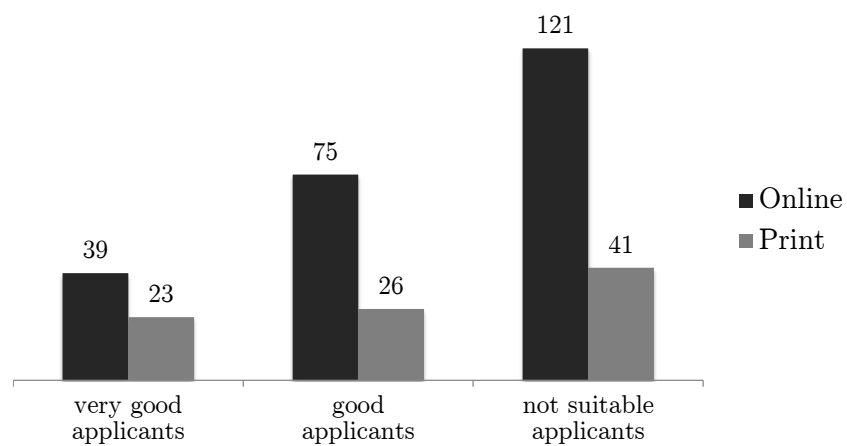
Source: Jobbörsen im Vergleich 2011, Fachhochschule Koblenz.  
Based on 1,500 radomly selected German job advertisements per website.

Figure 3.5  
Percentage of search results that match the search request on Monster.de



Source: Jobbörsen im Vergleich 2011, Fachhochschule Koblenz.  
Based on 86,023 results of the "quick search" function on Monster.de.

Figure 3.6  
Quality of applications in response to online and print job advertisements for "Head of Procurement"



Responses to job advertisements in 8 German newspapers (among others "Süddeutsche Zeitung" and "Der Tagesspiegel") and 8 German online job boards (among others monster.de and stellenanzeigen.de). Source: Medialeistungstest 2010, WESTPRESS GmbH & Co. KG.



Table 3.1  
All job changers and job changers who found their job through the Internet

Year	All	All Job Changers		Internet Job Seekers	
	Observations	Observations	% of total	Observations	% of changers
2000	24,576	2,102	8.55	21	1.00
2001	22,351	2,024	9.06	41	2.03
2002	23,892	1,901	7.96	64	3.37
2003	22,611	1,560	6.90	62	3.97
2004	22,019	1,378	6.26	75	5.44
2005	21,105	1,214	5.75	79	6.51
2006	22,665	1,380	6.09	89	6.45
2007	21,232	1,564	7.37	111	7.10
Total	180,451	13,123	7.27	542	4.13

*Notes:* The table shows the number of yearly observations in the full sample, the number and share of all job changers and the number and share of job changers who used the Internet to find their new job.

Table 3.2  
Sample means by job search method

	Internet	Newspaper	Friends	Private Agency	Job Center	Other	Back to former
Age	33.97 (9.16)	36.61 (9.95)	35.39 (10.94)	37.41 (10.44)	36.07 (11.45)	35.17 (10.27)	36.83 (10.59)
Male	0.58 (0.49)	0.40 (0.49)	0.46 (0.50)	0.62 (0.49)	0.53 (0.50)	0.48 (0.50)	0.34 (0.47)
Migrated	0.13 (0.34)	0.18 (0.38)	0.22 (0.41)	0.20 (0.40)	0.19 (0.39)	0.16 (0.37)	0.16 (0.37)
Education	4.13 (1.45)	3.66 (1.36)	3.39 (1.34)	3.96 (1.50)	3.28 (1.22)	3.79 (1.46)	3.63 (1.32)
Unemployed	0.43 (0.50)	0.30 (0.46)	0.30 (0.46)	0.37 (0.48)	0.69 (0.46)	0.26 (0.44)	0.23 (0.42)
<i>N</i>	542	1849	4116	181	1,426	2,185	1,315

*Notes:* The table shows the means of the respective subsample of job changer who used the Internet, newspapers, friends, private job agencies, job centers and other methods to find their new job. The last column shows the means of job changers who were reemployed by a former employer. Standard deviations in parentheses.

Table 3.3  
The association between online job search and several matching outcome variables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	skill use	work type	perspective	job security	benefits	work load	commute	work time
Found via Internet	0.0645*** (0.0214)	0.0798*** (0.0236)	0.0812*** (0.0202)	0.0455** (0.0212)	0.0687*** (0.0205)	0.0095 (0.0220)	-0.0070 (0.0235)	0.0150 (0.0237)
Male	0.0301*** (0.0106)	0.0682*** (0.0115)	0.0726*** (0.0103)	0.0615*** (0.0106)	0.0366*** (0.0104)	0.0044 (0.0108)	-0.0337*** (0.0112)	-0.0726*** (0.0115)
Age	-0.0074*** (0.0005)	-0.0072*** (0.0005)	-0.0096*** (0.0005)	-0.0049*** (0.0005)	-0.0030*** (0.0005)	-0.0016*** (0.0005)	-0.0010** (0.0005)	-0.0021*** (0.0005)
Migrated	-0.0304** (0.0130)	-0.0292** (0.0142)	-0.0453*** (0.0124)	-0.0271** (0.0127)	-0.0016 (0.0127)	-0.0053 (0.0131)	0.0070 (0.0140)	-0.0296** (0.0140)
Education	0.0327*** (0.0037)	0.0070* (0.0041)	0.0395*** (0.0036)	0.0104*** (0.0037)	0.0138*** (0.0036)	-0.0069* (0.0038)	-0.0035 (0.0040)	-0.0103** (0.0041)
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	9936	10065	9774	9796	9784	9994	9979	10022
Pseudo R2	0.098	0.076	0.121	0.079	0.073	0.058	0.050	0.067

Notes: The dependent variables take on the value 1 if the new job is evaluated better than the former one with respect to: the ability to use own skills, the type of work, the chances of promotion (perspective), the security against job loss, social benefits, work load, commute, and work time regulations. Individual-level covariates: unemployed, Internet at home and number of job changes. Average marginal effects of probit estimations with robust standard errors in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 3.4  
Regression-adjusted matching samples

	(1) skill use	(2) work type	(3) perspective	(4) job security
<b>Epanechnikov kernel</b>				
Found via Internet	0.0747*** (0.0198)	0.0743*** (0.0207)	0.0804*** (0.0189)	0.0377** (0.0187)
N	6430	6427	6253	6298
<b>5-nearest neighbor</b>				
Found via Internet	0.0605** (0.0241)	0.0745*** (0.0245)	0.0708*** (0.0230)	0.0285 (0.0230)
N	1715	1755	1660	1622
<b>1-nearest neighbor</b>				
Found via Internet	0.1002*** (0.0364)	0.0825** (0.0357)	0.0767** (0.0359)	0.0149 (0.0368)
N	593	665	590	581
Individual controls	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes
Industry fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes

*Notes:* Estimates from regression-adjusted propensity score matching. All listed covariates were also used for the estimation of the propensity score. Every shown coefficient comes from a separate regression with weights calculated from nearest neighbor, 5-nearest neighbor and Epanechnikov kernel matching algorithms. The dependent variables take on the value 1 if the new job is evaluated better than the former one with respect to: the ability to use own skills, the type of work, the chances of promotion (perspective) and security against job loss. Individual-level covariates: male, age, migrated, education, unemployed, Internet at home and number of job changes. Average marginal effects of probit estimations with robust standard errors in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 3.5  
Effect heterogeneity with the dependent variable "skill use"

	(1)	(2)	(3)	N
	Internet	Main effect	Interaction	
Re-entry into employment	0.0523* (0.0267)	-0.1031*** (0.0130)	0.1397* (0.0732)	9936
Female and children below age 16	0.0489* (0.0269)	-0.0991*** (0.0141)	0.1463** (0.0713)	9936
Rural area	0.0228 (0.0228)	-0.0285*** (0.0109)	0.1152** (0.0487)	9037
Unemployed during last 12 months	0.0850*** (0.0325)	-0.0313*** (0.0113)	-0.0284 (0.0474)	9936
Younger than 30	0.0833*** (0.0294)	0.0472*** (0.0160)	-0.0255 (0.0507)	9936
Tertiary education	0.0825** (0.0336)	0.0075 (0.0227)	-0.0183 (0.0477)	9936
Migration background	0.0580** (0.0261)	-0.0344** (0.0139)	0.1109 (0.0683)	9936

*Notes:* Dependent variable "skill use" takes on the value 1 if the new job is evaluated better than the former one with respect to the ability to use own skills. Every line represents average marginal effects of one probit estimation according to Ai and Norton (2003), with column (1) showing the effect of the "found by Internet" variable, column (2) showing the main effect of the variable in the respective row and column (3) showing the interaction effect of that variable with the "found by Internet" variable. Regressions include year, county and industry fixed effects and control for: male, age, migrated, education, unemployed, Internet at home. Robust standard errors in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 3.6  
Different job search channels compared to the employment office

	(1) skill use	(2) work type	(3) perspective	(4) job security
Found via Internet	0.0588** (0.0231)	0.0707*** (0.0256)	0.0676*** (0.0221)	0.0692*** (0.0231)
Found via friends	-0.0123 (0.0145)	0.0024 (0.0156)	-0.0089 (0.0142)	0.0508*** (0.0146)
Found via newspaper	0.0297* (0.0166)	0.0168 (0.0180)	0.0239 (0.0160)	0.0390** (0.0166)
Found via agency	-0.0420 (0.0384)	0.0418 (0.0411)	0.0080 (0.0363)	0.0229 (0.0375)
Found via other	0.0067 (0.0159)	0.0112 (0.0173)	0.0133 (0.0155)	0.0446*** (0.0161)
Individual controls	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes
Industry fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
N	9936	10065	9774	9796
Pseudo R2	0.100	0.083	0.129	0.086

*Notes:* The dependent variables take on the value 1 if the new job is evaluated better than the former one with respect to: the ability to use own skills, the type of work, the chances of promotion (perspective) and security against job loss. Individual-level covariates: male, age, migrated, education, unemployed, Internet at home and number of job changes. Average marginal effects of probit estimations with robust standard errors in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 3.7  
Reduced sample for comparison of Internet with newspaper and friends

	(1) skill use	(2) work type	(3) perspective	(4) job security
Found via Internet	0.0556** (0.0219)	0.0582** (0.0239)	0.0557*** (0.0209)	0.0125 (0.0221)
Male	0.0214 (0.0135)	0.0651*** (0.0146)	0.0834*** (0.0131)	0.0676*** (0.0136)
Age	-0.0073*** (0.0006)	-0.0079*** (0.0006)	-0.0096*** (0.0006)	-0.0052*** (0.0006)
Migrated	-0.0152 (0.0164)	-0.0280 (0.0176)	-0.0409*** (0.0154)	-0.0359** (0.0160)
Education	0.0316*** (0.0047)	0.0091* (0.0052)	0.0408*** (0.0046)	0.0128*** (0.0048)
Unemployed	-0.0382*** (0.0134)	-0.0504*** (0.0146)	-0.0477*** (0.0130)	-0.0364*** (0.0135)
Internet availability	0.0479*** (0.0137)	0.0266* (0.0148)	0.0114 (0.0134)	-0.0189 (0.0139)
Job changes	-0.0020 (0.0044)	-0.0178*** (0.0048)	-0.0128*** (0.0043)	-0.0150*** (0.0045)
County fixed effects	Yes	Yes	Yes	Yes
Industry fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
N	5630	5739	5528	5565
Pseudo R2	0.108	0.093	0.150	0.099

*Notes:* The dependent variables take on the value 1 if the new job is evaluated better than the former one with respect to: the ability to use own skills, the type of work, the chances of promotion (perspective) and security against job loss. The samples only include job changers who found their new job through friends, newspapers or the Internet. Average marginal effects of probit estimations with robust standard errors in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 3.8  
Internet job finders and personality traits

	(1) Optimism	(2) Satisfaction	(3) Originality	(4) Imagination
Found via internet	0.1105 (0.1002)	-0.1395 (0.0771)	-0.0109 (0.1808)	0.1806 (0.1905)
Male	-0.1049* (0.0481)	-0.0739* (0.0314)	0.2503** (0.0862)	-0.0662 (0.0985)
Age	0.0140*** (0.0023)	-0.0197*** (0.0016)	0.0013 (0.0042)	-0.0101* (0.0047)
Migrated	-0.0790 (0.0657)	0.0432 (0.0424)	-0.0259 (0.1179)	-0.1261 (0.1298)
Education	-0.0733*** (0.0181)	0.1149*** (0.0117)	0.0490 (0.0312)	0.0066 (0.0348)
Unemployed	0.1075* (0.0532)	-0.3825*** (0.0368)	-0.1721 (0.0974)	-0.2463* (0.1086)
Internet available	-0.0828 (0.0558)	0.0981** (0.0335)	0.2099* (0.1051)	-0.0237 (0.1141)
Job changes	-0.0097 (0.0179)	-0.0715*** (0.0118)	0.0398 (0.0316)	0.0485 (0.0374)
State fixed effects	Yes	Yes	Yes	Yes
N	1038	11581	1021	1019
R-sq	0.060	0.053	0.027	0.022

*Notes:* The dependent variables are self-assessed personality traits of the respondents: the degree of optimism when thinking about the future, overall life satisfaction, the degree of originality and the activeness of the own imagination. Ordinary least squares (OLS) estimations with robust standard errors in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.



Table 3.9  
Controlling for respondent's assessment of working time as a way to account for bias from systematic differences in job changer characteristics

	(1) skill use	(2) work type	(3) perspective	(4) job security
Internet	0.0619*** (0.0202)	0.0764*** (0.0222)	0.0736*** (0.0191)	0.0411** (0.0198)
Working time	0.0666*** (0.0091)	0.1764*** (0.0094)	0.0831*** (0.0087)	0.1528*** (0.0085)
Individual	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes
Industry fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
N	9810	10014	9738	9761
Pseudo R2	0.100	0.097	0.128	0.104

*Notes:* The dependent variables take on the value 1 if the new job is evaluated better than the former one with respect to: the ability to use own skills, the type of work, the chances of promotion (perspective), and the security against job loss. The "working time" variable takes on the value 1 if the new job is evaluated better than the former one with respects to the work hour regulations. As shown in Table 3.3, there is no significant correlation between finding a job through the Internet and improvements in working time. Individual-level covariates: male, age, migrated, education, unemployed, Internet at home and number of job changes. Average marginal effects of probit estimations with robust standard errors in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 3.10  
Controlling for whether job changers were actively searching for a job

	(1) skill use	(2) work type	(3) perspective	(4) job security
Found via Internet	0.0486** (0.0205)	0.0584** (0.0230)	0.0577*** (0.0195)	0.0316 (0.0203)
Active search	0.0636*** (0.0095)	0.0692*** (0.0103)	0.0574*** (0.0093)	0.0328*** (0.0095)
Covariates	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes
Industry fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
N	9880	10014	9724	9746
Pseudo R2	0.099	0.078	0.125	0.079

*Notes:* The dependent variables take on the value 1 if the new job is evaluated better than the former one with respect to: the ability to use own skills, the type of work, the chances of promotion (perspective) and the security against job loss. Individual-level covariates: male, age, migrated, education, unemployed, and Internet at home. Average marginal effects of probit estimations with robust standard errors in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

# Chapter 4

---

## PC use in Classrooms and Student Achievement\*

---

The use of computers in the classroom may be one of the most significant changes to teaching technology for decades, if not for centuries (e.g. Peterson, 2010). While the production of most goods has dramatically changed during the last few hundred years, the educational production technology in schools has experienced relatively little change. Indeed, the physical appearance of classrooms has not changed much since chalkboards were introduced in Prussian classrooms in the late 18th century (Konrad, 2007). However, the arrival of computers, tablets, and the Internet has challenged many traditional teaching practices and is generally seen as an opportunity for quality improvements.

Recently, there has been a big push in many countries to bring computers into classrooms. Although computers have been used in schools for decades, their use was often restricted to computer labs where students shared the PCs. Doubts as to the effectiveness of this practice (e.g. Oppenheimer, 2003) led to the idea of equipping every student in every classroom with a PC. But equipping every classroom with computers is expensive. Some school districts in the United States invest more than \$1 billion in classroom computers and corresponding infrastructure.<sup>1</sup> Indeed, President Obama has made technology in schools a priority of education policy and announced a multi-billion-dollar program to support the roll-out of technology in classrooms.<sup>2</sup> However,

---

\* This chapter was coauthored by Oliver Falck and Ludger Woessmann, University of Munich and Ifo Institute.

<sup>1</sup> The Los Angeles Unified School District plans to spend \$1.3 billion on iPads and Wi-Fi infrastructure. Source: <http://www.scpr.org/blogs/education/2014/02/11/15811/la-schools-wifi-networks-to-cost-about-800-million>

<sup>2</sup> Within the scope of the ConnectEd initiative, the Federal Communications Commission (FCC) will spend \$2 billion over the next two years to connect classrooms. Additionally, private companies, such as Microsoft and

empirical evidence on the effect of classroom computer use on student achievement is inconclusive at best, with many studies finding no or even negative effects.

We speculate that one reason for the mixed results of previous studies is that the effectiveness of PC use depends on the specific activities engaged in on the computer. In the case of home computers, it is often argued that the frequently found null effects are the result of positive effects caused by PC applications for schoolwork, combined with offsetting negative effects caused by applications for entertainment (Fairlie and Robinson, 2013). In the context of PC use in classrooms, the opportunity cost of time becomes more important. If teachers apply a new teaching method, they have to reduce the time that they spend on another method. The effect of the new teaching method on student achievement has therefore be interpreted in relation to the alternative method. We suggest that the positive effects of useful PC activities in the classroom are offset by the negative effects of activities that are inferior to alternative teaching methods. To date, there has been no systematic comparison of different PC use activities and their effectiveness for student achievement.

We use data on eighth-grade students from the international student achievement test TIMSS to show that three distinct PC activities have different effects on student achievement. Our identification strategy is based on different PC use intensities and test scores of students in math and science. This between-subject variation allows us to estimate within-student effects, holding unobserved school and student characteristics constant. By exploiting within-student between-subject variation, we also avoid any bias due to student selection into high PC usage schools or into high PC usage classrooms. To be a threat to our identification, students would need to select into schools or classrooms that have relatively higher PC usage in math compared to science and vice versa. Since school and class tracks are rarely separated by math and science and systematic differences in PC use between the two subjects are not very likely, we are confident that our identification method overcomes student selection bias.

Within-student identification strategies have recently gained traction in the economics of education literature. Starting with Dee (2005) who investigates the effect of teacher traits on student achievement, there have been studies using similar strategies to estimate the effects of same-gender teachers (Ammermüller and Dolton, 2006; Dee, 2007), teacher credentials and teacher subject knowledge (Clotfelter, Ladd and Vigdor, 2010; Metzler and Woessmann, 2012), instruction time (Lavy, 2010) and class size effects (Dee and West, 2011; Altinok and Kingdon, 2012). The work most closely related to our analysis uses within-student identification to investigate the effect

---

Apple, have committed more than \$1 billion to support the roll-out of new technologies into classrooms. Source: <http://www.whitehouse.gov/issues/education/k-12/connected>

of different teaching methods on student achievement (Aslam and Kingdon, 2011; Schwerdt and Wuppermann, 2011; Falch and Ronning, 2012). The majority of these studies assume that effects are the same across subjects. This common assumption of conventional fixed effects models will unlikely hold in the case of classroom PC use. Using PCs in science may have different effects on student achievement than using PCs in math.

Following Ashenfelter and Zimmerman (1997) and Metzler and Woessmann (2012), we apply a correlated random effects model that relaxes the assumption of equal PC use effects across subjects. In our model, we identify separate effects of PC use in math and science and are able to test whether effects are the same between subjects. For most specifications, we reject the assumption of equal effects and therefore decide against estimating conventional fixed effects models. Our correlated random effects model can be restricted such that the implied coefficients are numerically identical to first-difference estimates. Conventional fixed effects models thus can be considered to be nested within the more general correlated random effects models and there is no harm in estimating correlated random effects as opposed to fixed effects when the model is not overidentified (Ashenfelter and Zimmerman, 1997).

In TIMSS, we are able to differentiate between three PC activities: practicing skills and procedures, processing and analyzing data, and looking up ideas and information. While practicing skills and procedures is an activity that can also be performed without a PC, looking up ideas and information is quite PC specific. We find that the use of PCs for practicing skills and procedures has negative effects on math and science test scores. In science, the use of PCs for looking up ideas and information has a positive effect. Using the PC for processing and analyzing data has neither positive nor negative effects. The combination of the positive and negative effects for the different activities leads to an overall null effect of combined PC use in science and a small negative effect of combined PC use in math.

Our results show that one standard deviation higher PC use for looking up ideas and information increases student test scores in science by 3.3 percent of a standard deviation. For a class that currently does not have PCs available during science instruction, equipping students with PCs and letting them use the PCs to look up ideas and information almost every day, increases student test scores by about 11 percent of a standard deviation. However, equipping students with PCs and making them use the PCs for practicing skills and procedures almost every day, reduces test scores by about 9 percent of a standard deviation. The size of these effects is comparable to the impact of other teaching methods. For example, Schwerdt and Wuppermann (2011) find that

shifting from none to 100 percent of lecture style presentation increases student test scores by 10 percent of a standard deviation. Aslam and Kingdon (2011) show that explaining concepts and involving students in discussions raises test scores by 4.3 percent of a standard deviation. Falch and Ronning (2012) finds that assigning homework in all lessons compared to never assigning homework increases student test scores by about 4 percent of a standard deviation.

An important question for policy-makers is whether disadvantaged students benefit or suffer from increased PC use in the classroom. On the one hand, the use of PCs could make it easier for disadvantaged students to adjust the learning speed to their needs; on the other hand, the use of PCs requires more self-discipline and autonomy, which might be more difficult for disadvantaged students. We find that top-performing students are able to benefit slightly more from using a PC to look up ideas in science and suffer less from practicing skills at the PC in math. For low-performing and non-native students, looking up ideas in science does not have positive effects and is even detrimental when it comes to math. Students with lower socioeconomic background, approximated by a relatively low number of books at home, benefit from looking up ideas in science to a similar extent as the average student. We find no evidence that students who rarely use PCs at home benefit less from PC use in the classroom.

The cross-country setting of our study allows discovery of the country characteristics that favor effective use of PCs. We find that countries with a larger GDP per capita and a higher penetration of computers and broadband Internet use PCs in classrooms more frequently, but not necessarily more effectively. Specifically, in the richest or most populous countries, we find that students do not benefit at all from PC use in the classroom. Moreover, there is no systematic relationship of PC use effectiveness with the number of speakers of the predominant language in a country, which might increase the availability of digital teaching materials. For countries with high Internet usage and high broadband Internet penetration, we find positive effects of using the PC for looking up ideas.

The majority of this paper is focused on the eighth grade, a context in which we expect PC use to be most effective, but we also repeat our analysis for fourth-grade students. Although we have less information on the various PC use activities of fourth-grade students, we still find the general pattern that using the PC for practicing skills has negative effects on student test scores, whereas looking up ideas has positive effects. Many fourth-grade students are taught by the same teacher in math and science, which allows us to estimate within-student within-teacher effects, thus avoiding

bias from unobserved teacher characteristics or selection of teachers into classrooms.<sup>3</sup> There are very few differences between the estimates in the full fourth-grade sample and the within-teacher estimates in the same-teacher sample.

In addition to possibly enhancing student achievement in math and science, PC use in the classroom could have a whole range of other positive outcomes. Most importantly, classroom PC use could increase students' computer skills. Some argue that this is an even more important outcome than improving subject-specific student achievement.<sup>4</sup> We cannot make any claims regarding the effects on computer skills, but we believe that even if the primary objective of classroom PC use is to teach computer skills, it is nevertheless important to assess how subject-specific student achievement is affected. One way to interpret our results in this light is that there is no significant effect of combined PC use on test scores, at least in science, which is good news, assuming that some other goals are met by using the PC.

The remainder of the paper is organized as follows. Section 4.1 provides an overview of the previous literature. Section 4.2 introduces the TIMSS data. Section 4.3 describes our identification strategy and the correlated random effects model. Section 4.4 presents our main results. Section 4.5 and 4.6 describe student and country heterogeneities. Section 4.7 presents the results for our fourth-grade sample. Section 4.8 concludes.

## 4.1 Related Literature

This study contributes to the growing body of empirical work investigating the relationship between PC use and student achievement. Most of this work can be classified into one of three broad strands. The first strand consists of studies that analyze cross-sectional data to describe the correlation between computer use and educational outcomes. Most of these studies focus on computer use at home (e.g. Attewell and Battle, 1999; Fairlie, 2005; Schmitt and Wadsworth, 2006; Beltran, Das and Fairlie, 2010; Fiorini, 2010) and generally find a positive association between home computers and educational outcomes like cognitive skills, school enrollment or graduation rates. Fuchs and Woessmann (2004) use test data from the Programme for International Student Assessment (PISA) to show that after controlling extensively for student, school and family background characteristics, the

---

<sup>3</sup> One remaining concern is that teachers select into classrooms based on their relative PC use intensity these subjects.

<sup>4</sup> For example, Karen Cator, Director of the Office of Educational Technology in the U.S. Department of Education, believes that standardized test scores are an inadequate measure of the value of technology in schools. Source: <http://www.nytimes.com/2011/09/04/technology/technology-in-schools-faces-questions-on-value.html>

initial positive association between computer use at home and student achievement becomes negative. The authors find no effect on student performance from the use of computers at school.

The second strand of literature uses exogenous variation from government programs to identify the effect of computer use at school. Angrist and Lavy (2002) exploit an Israeli government program that provided more than 50,000 computers to schools. The authors do not find any positive effects on student achievement. Goolsbee and Guryan (2006) use variation from the E-Rate program in the United States, which provides up to 2.25 billion dollar per year for better computers and Internet connections at schools and libraries. Although the program significantly increased ICT investment by schools, the authors do not find positive effects on student performance. Machin, McNally and Silva (2007) exploit a change in strategy by the U.K. government that resulted in some British primary schools increasing their ICT investment. They find that more ICT funding had a positive effect on student achievement for English and science, but not for math. The studies by Leuven and Lindahl (2007) and Malamud and Pop-Eleches (2011) use regression discontinuity designs for identification. Leuven and Lindahl use a program targeted at Dutch primary schools having more than 70 percent disadvantaged students and find no positive effects of extra funding for computers and software. Malamud and Pop-Eleches analyze a voucher program for students from low-income families in Romania and find that computers at home do not improve student achievement.

The third strand of literature consists of randomized controlled field experiments on PC use in schools. Barrera-Osorio and Linden (2009) randomly divide more than 100 schools in Colombia into a treatment and control group. The treatment group received computers for students and training for teachers, the control group did not. The authors find that after two years, student achievement in the treatment group did not exceed achievement in the control group. Cristia et al. (2012) evaluate the effect of the One Laptop per Child Program on students in 319 primary schools in Peru. They find no positive effects on student enrollment rates or test scores. Several field experiments focus on the use of specific e-learning software. Rouse and Krueger (2004) show that use of the software "Fast ForWord" can help low-performing students do better on computer-based tests, although there was no effect on other standardized language tests. Banerjee et al. (2007) conduct a field experiment with about 6,000 Indian students and find that low-performing students do better on mathematics tests when they regularly use specific training software. Banerjee et al. (2007) conducts an experiment in India in which students receive computer-based training either in addition to or instead of traditional class lectures. Students who use the computer instead of having normal lessons perform worse than others; students who use the computer in addition to normal lessons perform better than the control group. Fairlie and Robinson (2013) is one of the few randomized experiments



that analyzes the use of home computers on student achievement. The authors find no effects on grades, test scores, attendance, or other educational outcomes of students.

The majority of the extant literature has limitations either with regard to a causal interpretation of the effects or with external validity. Results from the first strand of literature, which uses cross-sectional data, are likely to suffer from omitted variable bias. It is obvious that the use of computers at school is closely related to school and teacher characteristics that are usually unobservable. Furthermore, this kind of study does not overcome the problem that students with unobserved ability and family background select into schools and classrooms. The second strand of literature avoids this bias by exploiting exogenous variation from government programs. However, these programs merely increase the availability and not necessarily the usage of PCs. Additionally, the students affected by the exogenous variation often belong to a specific group, for example, low-income students. The third strand of literature uses experiments in order to provide a clean identification of causal effects. Since many of these experiments are conducted in developing countries, their external validity for the use of PCs in more developed countries is questionable. Experiments with specific e-learning software suffer from the same problem of external validity, as software used in one country is often very different from e-learning software used in another country.

Our study makes three contributions to the literature. First, we provide results that allow for a causal interpretation of PC use effects on achievement across countries. Second, since our study encompasses a wide variety of countries and different PC use applications, we are able to explore heterogeneities and the external validity of our results. Third, and most importantly, we are the first to systematically analyze the effects of three different PC use activities, namely, practicing skills, processing data, and looking up ideas. Opening this black box of PC use, we provide insight into the relative effectiveness of different PC activities and, at the same time, offer an explanation of the inconclusive results found in previous studies.

## 4.2 Data

To estimate the within-student between-subject effect of classroom PC use on student achievement, we use data from the 2011 wave of the Trends in International Math and Science Study (TIMSS). The TIMSS data have several features that make them especially interesting for our analysis. Unlike other studies on the effect of PC use on achievement that use either very specific country settings

(e.g., Banerjee et al., 2007 and Barrera-Osorio and Linden, 2009 who focus on developing countries) or studies that use programs targeted at a specific group of students (e.g., Leuven and Lindahl, 2007 and Malamud and Pop-Eleches, 2011 who focus on disadvantaged students), TIMSS data include students from a variety of countries and from all sorts of backgrounds. In our eighth-grade sample, we have students from 33 educational systems, which allows us to analyze whether the effectiveness of PC use is correlated with observable country characteristics like GDP per capita or share of the population using the Internet. To ensure a representative sample of students in each country, TIMSS employs a two-stage sampling design. In the first step, out of all schools in a country, a certain number are sampled with sampling probabilities proportional to school size. In the second step, one or more classrooms from those schools are randomly sampled. This sampling design allows a representative analysis of schools, teachers, classrooms, and students in every country.<sup>5</sup>

TIMSS tests students in the fourth and eighth grades. Students in the fourth grade are usually around 10 years old; students in the eighth grade are around age 14. Not all countries participate in both the fourth-grade and the eighth-grade test and not all countries participate in both subjects, math and science. For most of this paper, we focus on the eighth-grade sample of TIMSS 2011, for two reasons. The first reason is that it is often hypothesized that the use of PCs in the classroom requires a certain degree of maturity and thus classroom PC use might be more effective in post-elementary school settings. It is also likely that eighth-grade students use PCs in a more systematic way than do fourth-graders. Our second reason for focusing on eighth-grade students is that we have more information about PC activities in the eighth-grade sample. TIMSS asks eighth-grade teachers more questions about PC use, likely because at that stage of schooling students can perform more specific tasks on the PC, for example, processing data from experiments in science classes. In the fourth grade, PC use is usually more playful and less task oriented.

Since the main focus of this study is to look closely at the types of PC activities and their effectiveness, we prefer to work with the eighth-grade sample. However, for purposes of identification, the fourth-grade sample has one advantage: many students are taught math and science by the same teacher. By restricting the fourth-grade sample to students with the same teacher in both subjects, we can avoid potential bias from teacher selection into classrooms. Therefore, we repeat most of our analysis of eighth-grade students with fourth-grade students and present the results in Section 4.7. Comparing the results from the whole fourth-grade sample with the results from the fourth-grade same-teacher sample, it seems that teacher selection is not a major concern. Hence, we choose the

---

<sup>5</sup> Students with disabilities and students who are unable to read or speak the national TIMSS test language are not sampled in TIMSS.

greater information about PC activities in the eighth-grade sample over the identification advantage in the fourth-grade sample and from this point on mainly discuss the data, identification strategies, and results for eighth-grade students.

TIMSS also meets an important data requirement for our identification strategy, namely, the availability of comparable test scores from at least two different subjects. In TIMSS, students are tested in math and science and for ease of interpretation we scale test scores with a mean of 0 and a standard deviation of 1. TIMSS also offers sufficient variation of PC use intensity between subjects. Unlike other student achievement tests, such as the Programme for International Student Assessment (PISA), TIMSS not only includes student and school questionnaires, but also extensive teacher questionnaires. Teachers answer a wide set of questions regarding their personal characteristics, such as sex, age, and education, and are also questioned extensively about their teaching methods. This rich information about teacher characteristics and teaching methods allows us to mitigate potential bias due to the correlation of teacher traits with the use of PCs in the classroom. Assume, for example, that young teachers are both more likely to use PCs and more motivated to teach. If we simply estimate the effect of PC use on student achievement without taking teacher age into account, our estimations will suffer from omitted variable bias. The same would be true if teachers who use PCs in the classroom generally assign more homework. It is thus essential to control extensively for teacher characteristics and teaching methods, which is possible thanks to the rich TIMSS data on teachers.

Among the information on teaching methods, we are most interested in data on PC use in the classroom. Here, the TIMSS data again facilitate our estimation strategy, including as they do, and unlike any other student achievement test of which we are aware, detailed and yet comparable information on PC use in two different subjects. In TIMSS, teachers are asked: "Do the students in this class have computer(s) available to use during the lessons?" Given PC availability, teachers are then asked if and how they use the PCs. The question reads: "How often do you have the students do the following PC activities during the lessons?" The three PC activities are: "practice skills and procedures," "look up ideas and information," and "process and analyze data." The possible answers for each activity are: "every or almost every day," "once or twice a week," "once or twice a month," and "never or almost never." Note that for the fourth-grade sample, there is only information on two PC activities: "practice skills and procedures" and "look up ideas and information."

The three PC activities (henceforth, "processing data," "practicing skills," and "looking up ideas") strike a good balance between concreteness and abstractness. They are concrete enough

to get a sense of what students actually do with the PC, while still being applicable to both math and science lessons. This enables us to look more closely at the effectiveness of different PC uses in the classrooms while ensuring a clean identification of within-student between-subject effects. Most previous studies that go beyond the mere availability of PCs and try to analyze the effect of different PC activities use much more specific PC activities. For example, Rouse and Krueger (2004) analyze the effect of the "Fast ForWord" software program. The problem with studying very specific applications is that the findings have limited external validity. For example, it is questionable whether the results hold for follow-up versions of the same software, similar software offered by different developers, software used in different countries and languages, or software that applies similar principles to different subjects. In our case, students might not all use the same software if the teacher indicates that PCs are used to practice skills, but the external validity of our results should be relatively high.

The intensity of the three PC activities is measured from "never or almost never" to "every or almost every day." The most obvious way to arrive at a single measure of usage intensity would be to create a linear variable that goes from 0 for the lowest category to 4 for the highest category. However, doing so would have an important drawback in terms of interpretation. In most specifications, we include all three activities at the same time based on the idea that although the overall effect of all activities combined might be close to zero, as other studies have shown, some specific activities could have positive effects and others negative effects. However, the distribution of usage intensity is not the same for all activities. Moving up one category in the intensity of looking up ideas could be a much larger step than moving up one category in the intensity of practicing skills. To overcome this problem we create standardized usage intensity scores with a mean of 0 and a standard deviation of 1. The coefficients for all activities thus are comparable and can be interpreted as changes in a standard deviation.

Standardizing PC use intensity has the additional advantage that the intensities easily can be added up to a single measure of combined PC use. Although we will be most interested in the coefficients of each separate PC activity, it is useful to have a combined measure of all PC use activities. This will not only allow us to analyze the overall effect of PC use, but will also serve as a benchmark for other studies of PC use in the classroom at a more aggregated level. We create the combined PC use measure by summing all standardized measures for every PC activity and then standardize it again to mean 0. Note that due to this design, practicing skills on a daily basis does not necessarily have the same importance for the combined index as looking up ideas on a daily basis. However, it seems more sensible to give extraordinary high usage intensities of one activity

the same importance as extraordinary high intensities of another activity. We use different designs of the standardized scores as robustness checks.

Table 4.1 shows descriptive statistics for the PC variables in math and science. The first two rows show that 34 percent of all students have PCs available in math classrooms and that 43 percent have PCs available in science classrooms. The binary variables in the next six rows are 1 if the respective PC activity is performed at least "once or twice a month." For all activities, the share of students who regularly perform the activity in math is smaller than the share of students performing the activity in science. The largest difference is in looking up ideas, which is engaged in by 17 percent of students in math and by almost twice as many students in science. The share of students performing any of the three PC activities in math does not vary much between activities. Note that on the micro level, PCs are not generally used more frequently in math than in science.

Rows 9 to 14 of Table 4.1 show PC usage intensity by activity with linearly coded categories from 1, the lowest usage intensity, to 4, the highest intensity. The average usage intensities for most variables lie between 1.6 and 1.9, which means somewhere between "once or twice a month" and "once or twice a week." Looking up ideas in science is an exception, having an average usage intensity of 2.16, which makes it the most frequently engaged in PC activity. Note, however, that the variable means in Rows 9 to 14 only include students who have PCs available in the classroom. Therefore, the question answered by the mean is: Given that students have PCs available for usage in class, how often do they use them for processing data, practicing skills, and looking up ideas?

Although it is interesting to analyze usage intensity given PC availability, in most specifications we choose to include in our regressions students without PCs and control for PC availability. The reason for this choice is that the availability of PCs could be endogenous. The availability question in TIMSS reads: "Do the students in this class have computer(s) available to use during the lessons?" It is likely that teachers who choose not to use PCs for instruction will answer the question in the negative. The question, therefore, captures not only some potentially exogenous IT infrastructure, but also includes variation that is due to teacher choice. For this reason, we take students without available PCs into account by using the same linear variable as in Rows 9 to 14 but with the difference that students without PCs take the value 0. Due to the large number of zeros, the means in Rows 15-20 are considerably lower than those in Rows 9-14. Nevertheless, the pattern remains the same, that is, higher usage intensity in science compared to math and the highest intensity for looking up ideas in science. The variables in Rows 15-20 are the basis for the standardized variables used in most of our specifications.

Our cross-country setting allows analysis of country heterogeneities. We investigate whether certain country characteristics hinder or support the effectiveness of PC use in school. To this end, we use data from different sources and merge them with our TIMSS dataset. For general country characteristics, like GDP per capita and population, we use 2010 data from the World Bank. We speculate that the effectiveness of PC use could be correlated with the number of people who speak the country's predominant language. The rationale is that there is generally less digital learning material and software available for languages with few speakers. We use information on the number of speakers of every language from the Ethnologue 2013, which is considered the most comprehensive and accessible language catalog. We also hypothesize that the effectiveness of PC use could be correlated to a country's telecommunication infrastructure. Therefore, we use data from the World Telecommunication/ICT Indicators Database, which is published by the ITU (International Telecommunication Union), a U.N. agency for information and communication technologies. The database includes over 150 ICT indicators, which are collected from the national telecommunication ministries, regulatory authorities, and national statistical offices of more than 200 economies. In this paper, we mainly use the percentage of individuals using the Internet and the number of fixed broadband Internet subscriptions per 100 inhabitants.

### **4.3 Empirical Framework**

We identify the causal effect of classroom PC use on student achievement by exploiting within-student between-subject variation. A number of studies use the same strategy to avoid bias from omitted variables and selection when estimating the effect of non-subject-specific teacher characteristics on achievement (e.g., Dee, 2005; Ammermüller and Dolton, 2006; Clotfelter, Ladd and Vigdor, 2010; Metzler and Woessmann, 2012). Many comparable studies take a fixed effects approach implemented in the form of first-difference estimations. When we estimate the within-students between-subject effect in first differences, we implicitly assume that the effect is the same across subjects. While this assumption might be realistic in some applications, we doubt that this is the case for PC use.

There are several reasons why the use of PCs could have different effects in math and in science. First, it seems very likely that PCs are used differently in the two subjects. Our experience from visiting a leading German school in which computers are used in all subjects is that in every subject teachers and students use PCs slightly differently. In math, for example, many teachers in Germany use the GeoGebra software, which allows students to graphically explore different fields

of geometry, algebra, calculus, and statistics. In science, students are more likely to use the Internet for research on specific topics and to work on projects with other students. Second, teachers might find it easier to integrate PCs into teaching one subject compared to another subject. For example, it might be relatively easy for a teacher to use GeoGebra in math because the software is free and the teaching materials are readily available. For other subjects it could be more difficult to find suitable software packages that are both free and accompanied by useful materials. Third, some subjects might simply be better suited for PC use than others. For example, maybe using PCs several times a week for learning algebra works well, but is less useful for chemistry; maybe PCs are generally suitable for science instruction, but not so good for learning math.

Differences in the effect of PC use between math and science would lead to a bias of the conventional fixed effects estimator. Therefore, following Ashenfelter and Zimmerman (1997) and Metzler and Woessmann (2012), we adopt a correlated random effects model that allows us to relax the assumptions of the first-difference estimator. We begin by defining an educational production function for math and science with an explicit focus on PC use:

$$y_{mi} = \beta_m pc_{mt} + \gamma T_{mt} + \delta X_{mi} + \alpha Z_i + \mu_i + \tau_{mt} + \varepsilon_{mi} \quad (4.1)$$

$$y_{si} = \beta_s pc_{st} + \gamma T_{st} + \delta X_{si} + \alpha Z_i + \mu_i + \tau_{st} + \varepsilon_{si} \quad (4.2)$$

where  $y_{mi}$  and  $y_{si}$  are the test scores of student  $i$  in math and science. The educational input we are most interested in is the intensity of PC use in the classroom, denoted by  $pc_t$ , which is determined by teacher  $t$  for each subject individually. We also include a vector of teacher characteristics and teaching methods  $T_t$  that includes, for example, the age and education of the teacher and the amount of homework usually assigned by the teacher. The teacher characteristics in a narrower sense (like teacher age and education) vary between equation (1) and (2) only when the student is taught by different teachers in the two subjects. The teaching method controls even vary between subjects when the student is instructed by the same teacher in both math and science. We also control for a large set of student, family, school and classroom characteristics that are both subject-specific ( $X_i$ ) and non-subject-specific ( $Z_i$ ). The error term is divided into a student-specific component  $\mu_i$ , a teacher-specific component  $\tau_t$ , and a subject-specific component  $\varepsilon_i$ . All estimations are weighted.

There are many reasons to believe that the student-specific error term  $\mu_i$  is correlated with classroom PC use. If this is the case, Equations (1) and (2) will not appropriately identify the effect of PC use on student achievement. It is obvious, for example, that for there to be classroom PC use there must be PCs in the classroom. The reasons why some schools have PCs and others do not are anything but exogenous. Schools in economically stronger areas might find it easier to attract money from parents or local firms for new computer equipment. Students in these areas are also likely to have a higher socioeconomic status, which is positively correlated with student performance. Within schools, it might be that teachers use PCs only with the best performing classes because they do not want to take chances with low-performing students. In some cases, PC classes are opt-in, which means that the parents can choose whether their child should be in a "regular" class or in a class where PCs are used extensively. In these cases, we would have a strong correlation between unobserved student traits and classroom PC use. Therefore, following Chamberlain (1982), we model the student-specific error component which varies across students but is the same for one student between math and science:

$$\mu_i = \eta_m pc_{mt} + \eta_s pc_{st} + \theta_m T_{mt} + \theta_s T_{st} + \phi X_{mi} + \phi X_{si} + \chi Z_i + \omega_i \quad (4.3)$$

where the residual  $\omega_i$  is uncorrelated with the observed variables in  $pc_t$ ,  $T_t$ ,  $X_i$  and  $Z_i$ . Note that the parameters  $\eta$  and  $\theta$  may take different values in math and science. Substituting Equation (3) into Equation (1) and (2) we obtain:

$$y_{mi} = (\beta_m + \eta_m) pc_{mt} + \eta_s pc_{st} + (\gamma + \theta_m) T_{mt} + \theta_s T_{st} \\ + (\delta + \phi) X_{mi} + \delta X_{si} + (\alpha + \chi) Z_i + \tau_{mt} + \varepsilon_{mi} \quad (4.4)$$

$$y_{si} = (\beta_s + \eta_s) pc_{st} + \eta_m pc_{mt} + (\gamma + \theta_s) T_{st} + \theta_m T_{mt} \\ + (\delta + \phi) X_{si} + \delta X_{mi} + (\alpha + \chi) Z_i + \tau_{st} + \varepsilon_{si} \quad (4.5)$$

Note that all subject-specific controls enter both the equation for math and the equation for science. Equations (4) and (5) comprise a correlated random effects model that can be estimated straightforwardly. Compared to a conventional fixed effects estimation, this model has at least three advantages. First, and most importantly, we do not need to assume that the effect of computer use



is the same for each subject and we can even test whether this assumption would hold, namely if  $\beta_m = \beta_s$ . Second, despite having within-student estimations, we can estimate the effect sizes  $\beta_m$  and  $\beta_s$  separately and thereby evaluate for which subject it is efficient to use computers and for which subject it is not. Third, we can estimate the bias that OLS estimations of Equations (1) and (2) would suffer. Specifically,  $\eta_m$  and  $\eta_s$  measure how much the  $\beta$  coefficients in Equations (1) and (2) would be biased due to omitted country, school, family and student characteristics that influence student test scores. For example,  $\eta_m$  and  $\eta_s$  should be positive if students with higher socio-economic status are more likely to use computers in the classroom. We can also test whether  $\eta_m = \eta_s$ . If we are not able to reject this test, we can conclude that the observables we control for parameterize the model reasonably well.

The conventional fixed effects model is actually nested within the correlated random effects model in Equations (4) and (5). To better see this relationship, we assume  $\beta_m = \beta_s$  and  $\eta_m = \eta_s$ , which are the two implicit assumptions of conventional fixed effects models, and build the difference between (4) and (5):

$$y_{mi} - y_{si} = \beta(p_{cmt} - p_{cst}) + \gamma(T_{mt} - T_{st}) + \delta(X_{mi} - X_{si}) + \tau_{mt} - \tau_{st} + \varepsilon_{mi} - \varepsilon_{si} \quad (4.6)$$

Equation (6) is equivalent to the conventional first-difference estimator with student fixed effects. Note that fixed effects on the country, school and family level are all nested within student fixed effects and thus do not need to be included separately. Due to the relationship between the correlated random effects model and the fixed effects model, it is always possible to estimate fixed effects within the correlated random effects framework by simply restricting the  $\beta$  and  $\eta$  coefficients to be equal across subjects. This implies that there is no harm in estimating correlated random effects instead of fixed effects, while estimating correlated random effects has the above-mentioned advantages of testing the fixed effects assumptions and identifying the effects for each subject separately.

One additional concern is that unobserved family and student characteristics that are absorbed by the student fixed effects matter more for PC use in one subject than in the other. For example, a school might have computers available in most science classrooms, but in only a few math classrooms. While the socioeconomic status of a student would in this case have little effect on PC

use in science, its influence on PC use in math would be much stronger. Although these scenarios do not seem very likely, they could lead to an exaggeration of the difference between  $\beta_m$  and  $\beta_s$ . To discover whether this is indeed a problem, we can denote the importance of unobserved student variables in math relative to science with  $\alpha$  and write:

$$y_{mi} = (\beta_m + \alpha\eta_m)pc_{mt} + \alpha\eta_s pc_{st} + (\gamma + \theta_m)T_{mt} + \theta_s T_{st} \\ + (\delta + \phi)X_{mi} + \delta X_{si} + (\alpha + \chi)Z_i + \tau_{mt} + \varepsilon_{mi} \quad (4.7)$$

$$y_{si} = (\beta_s + \eta_s)pc_{st} + \eta_m pc_{mt} + (\gamma + \theta_s)T_{st} + \theta_m T_{mt} \\ + (\delta + \phi)X_{si} + \delta X_{mi} + (\alpha + \chi)Z_i + \tau_{st} + \varepsilon_{si} \quad (4.8)$$

To estimate (6), we need to impose one additional restriction, for example  $\eta_m = \eta_s$ , so that all parameters are exactly identified.

For our within-student between-subject estimations, it is important that there is enough variation in the intensity of computer use between subjects. On the country level, there is a great deal of variation between subjects. In some countries, the vast majority of students use PCs for science and a much smaller proportion use them for math; in other countries it is the other way around. More importantly, there is also a lot of variation on the student level. If students used PCs either in both subjects or not at all, we would not have the necessary variation to identify a within-student between-subject effect. Table A4.1 shows the number and share of observations with different PC use intensities by country. In Australia, for example, more than 80 percent of students have different PC use intensities between math and science for our three PC use categories of "processing data," "practicing skills," and "looking up ideas." Iran has the lowest within-student variation, with less than one-third of students using PCs to a different degree in math and science. Note that even in some countries with low overall PC use intensity, such as Japan, we observe different intensities between subjects for around 50 percent of the students. Although in most countries PCs are used more often in science than in math, on the student level there is a surprisingly high proportion of students who use PCs more often in math.

## 4.4 Results

As a benchmark for previous studies, we start with conventional OLS estimations before turning to our correlated random effects model. Table 4.2 shows the results of least-squares estimations in the eighth-grade TIMSS sample with different sets of control variables. In Table 4.2, as well as in most of the following tables, we show estimates of two different estimations. The first estimation, shown in the Row 1, includes the combined index variable that subsumes PC usage for processing data, practicing skills, and looking up ideas. In the second estimation, shown in Rows 2-4, the three PC activities enter the regression separately but at the same time. The combined PC use estimation serves as a benchmark for the majority of previous studies that do not analyze the effect of PC use on the level of specific PC activities. We expect to find small positive or negative effects of the combined PC usage measure, comparable to the effects that other studies have found. In the estimations with separate PC activity measures, we expect that some estimates will be positive, others negative.

Specifications 1, 2 and 3 of Table 4.2 each show separate regressions with student test scores in math and science as the respective dependent variable. Consistent with Fuchs and Woessmann (2004) and many other studies, there is a positive bivariate relationship between PC use and student achievement. Model 1 shows that without any controls, an increase of PC usage by one standard deviation is associated with 11 percent of a standard deviation higher test scores in math. For science, the association is even stronger. As previous studies show, the positive association between test scores and combined PC usage vanishes as more controls are included in the model. In Model 3, which has full sets of student, teacher, and teaching method controls, the association becomes very small, statistically insignificant, and even turns negative for math.

Looking at the lower panel of Table 4.2, where the estimations with separate PC activity variables are shown, we see that there is a great deal of heterogeneity between activities. The coefficients for processing data are generally positive; the coefficients for practicing skills are often negative after controlling for student and teacher characteristics. In the specification with full controls, estimates for looking up ideas are low and not significant. There is also quite a difference between the association of the PC activity variables with student test scores in math and their association with student test scores in science. Practicing skills, for example, has no clear association with test scores in science but a highly significant negative association with test scores in math. The different

coefficients for math and science raise strong doubts that the conventional fixed effects assumption of equal effects across subjects holds in the case of PC use in math and science.

In specification 4 of Table 4.2, we repeat specification 3 but estimate it as a seemingly unrelated regression (SUR). It is highly likely that the errors of the estimation with test scores in math are correlated with the errors of the estimation with test scores in science as dependent variable. It is worth noting in this context that the control variables of the math and science regressions are not entirely the same. Student controls do not vary between subjects, but teacher controls are the same between math and science only if a student is taught both subjects by the same teacher. Teaching method controls are always subject-specific. Due to the likely correlation of errors between the regressions for math and science, we estimate SUR for the remainder of the paper. We also cluster on the classroom level in all our estimations. Since in specification 4, the sample size is slightly reduced due to the further restriction that all covariates of both subjects have to be non-missing, the coefficients change marginally. More importantly, however, the standard errors are smaller in SUR compared to OLS. As a result, the negative coefficient for combined PC use becomes highly significant in math.

Table 4.3 presents our correlated random effects model of Equations (4) and (5). Following equations (4) and (5), the effects of PC use in math shown as "Implied  $\beta$ " in the math columns of Table 4.3 are calculated as the difference between the coefficient of PC use in math in Equation (4) and the coefficient of PC use in math in Equation (5). The implied  $\beta$  for science is the coefficient of PC use in science in Equation (5) minus the coefficient of PC use in science in Equation (4). Since the simple  $\beta$  and  $\eta$  coefficients of Equations (4) and (5) have no straightforward interpretation, they are not shown in Table 4.3. After calculating the implied  $\beta$  for math and science, we test whether they are significantly different from zero. The  $\chi^2$  test statistics are shown in brackets below the implied  $\beta$  coefficients and the  $Prob > \chi^2$  values are reflected by the stars behind the implied  $\beta$  coefficients, following the conventional significance levels of 0.01, 0.05, and 0.1.

Below every implied  $\beta$  coefficient and the corresponding  $\chi^2$  statistic, Table 4.3 shows the difference between the implied  $\beta$ s and the  $\eta$ s in math and science. We test whether the coefficients are significantly different between subjects. Due to space limitations, we do not show the  $\chi^2$  test statistic but we denote the significance levels by stars behind the shown difference between the coefficients. Whenever we fail to reject that the implied  $\eta$ s are different from each other ( $\eta_{math} - \eta_{science} = 0$ ), we could also estimate a restricted model that assumes  $\eta_{math} = \eta_{science}$ . When at the same time, we cannot reject that the implied  $\beta$ s of the two subjects are different from

each other, we could in fact estimate a conventional first-difference model as presented in Equation (6). However, in most specifications at least one variable of interest has implied  $\beta$  or  $\eta$  coefficients that are significantly different from each other. Since the restriction of  $\eta_{math} = \eta_{science}$  for a single variable results in only minor efficiency gains for our estimation, we decide against applying such restrictions. The differences between the  $\beta$  coefficients strongly indicate that the  $\beta_{math} = \beta_{science}$  assumption of conventional fixed effects models is violated and advise us against estimating such models for our application.

Table 4.3 shows that the combined effect of PC use is zero when some basic controls are included in the model. For math, the combined effect even becomes significantly negative in the specification with full sets of controls. It is striking how similar the coefficients of combined PC usage in specification 3 of Table 4.3 are to the comparable specification 4 in Table 4.2. One could conclude that, at least for the estimations with combined PC usage, the control variables account quite well for school, family and student characteristics. The zero and negative effects on test scores are consistent with many previous studies that look at the overall effect of PC use in schools. Notable examples of studies with rigorous identification strategies include Angrist and Lavy (2002), who find no effect, and Leuven and Lindahl (2007), who find negative effects. There are some studies that find a similar pattern of effects between math and science. For example, Machin, McNally and Silva (2007) find a slightly positive effect of ICT investment in schools on student test scores in science but fail to find any impact on test scores in math. Similarly, Cristia et al. (2012) find positive effects of PC use on several skill dimensions, but not on math test scores. Therefore our estimates of the combined PC use effect are broadly consistent with previous finding.

The main contribution of our paper lies in the analysis of PC use activities, rather than in the analysis of the combined PC use for all kinds of activities. Therefore, we are most interested in the lower panel of Table 4.3 which shows the effects of using a PC for processing data, practicing skills, and looking up ideas. The effects of using a PC for processing data are all very close to zero. In our main specification with all controls, the effects in math and science are both smaller than 1 percent of a standard deviation. For the other two activities, we find some statistically significant effects. In specification 1, which includes only basic controls, the effects of using a PC for practicing skills are negative in both subjects. The effects of using a PC for looking up ideas are of a similar magnitude, but positive. Having one activity with an effect of zero, one activity with negative effects, and one activity with positive effects results in the null effect we find for the combined PC usage.

Interestingly, for processing data and practicing skills, the effects on test scores are not significantly different between math and science. For looking up ideas, the difference between the effects is also small when only basic controls are included. However, the difference increases with additional controls. The effect of looking up ideas in science remains positive and statistically significant, but the effect of looking up ideas in math is almost halved by controlling for teacher characteristics and reduced further to 0.5 percent of a standard deviation when all controls are included. This difference in the effect of looking up ideas also drives the difference in the combined PC use effect between math and science.

One could argue that controlling for teaching characteristics is not necessarily desirable because, for example, the time spent on teacher-centric lecturing could also be an outcome variable of PC use. The same could be said about other teaching method controls, like the number of examinations and the amount of homework assigned. However, Table 4.3 shows that there is little difference between specification 2, which only controls for teacher characteristics like sex, age, and education, and specification 3, which includes teaching method controls. There are also plausible explanations for why some of the PC use coefficients become smaller after teaching methods are controlled for. For example, teachers who use PCs in the classroom might be generally more innovative than teachers who do not use PCs. It could also be considered more innovative to discuss a student's progress with his or her parents, to praise students regularly for their effort, and to frequently summarize what students should have learned. By not including these controls, we potentially overestimate the effect of PC use on student achievement. We thus use specification 3 as our baseline model in the remainder of the paper.

## 4.5 Student Heterogeneities

From the results in Table 4.3 we learned that using a PC for practicing skills has a negative effect on student achievement, whereas using it to look up ideas has a positive effect on student achievement in science. In the next step, we want to discover whether some students benefit more than others by engaging in certain PC activities. There are many competing hypotheses about who benefits most from classroom PC use. Some studies argue that the weakest students benefit the most because they can repeat learning materials as often as they want on a PC. Learning on a PC could also be better calibrated to a student's capabilities, thus avoiding the creation of knowledge gaps in the first place. There are several empirical studies that look specifically at students with learning difficulties. A good example is the above-mentioned study by Rouse and Krueger (2004) who analyze the effect of

the language-learning software "Fast ForWord". Rouse and Krueger restrict their sample to students who scored in the bottom 20 percent on the state's standardized reading test. The study finds small positive effects from using "Fast ForWord" on reading skills as measured by an assessment test that is sold by the same company that sells Fast ForWord; however, the authors find no effects on the test scores from other standardized language tests.

A competing hypothesis is that the best students benefit the most from PC activities in the classroom. The arguments in support of this idea often cite the autonomy and responsibility that students have when working with a PC. While it might be true that, theoretically, the weakest students can go over learning materials as often as they want or need to on a PC, the question is whether they really do. Very often, the use of PCs in school entails students being proactive and engaging with different learning tools or the Internet by themselves. The students who are most proactive, however, are usually the strongest students, not the weakest. Using a PC to look up ideas requires students to work without immediate support from the teacher and resist, for example, the temptation to simply surf the Internet for fun, which can be more challenging for weaker students than for stronger students. Another argument in this vein is that there are skill complementarities in PC use. To use a PC to their advantage, students might need skills such as creativity and critical thinking. If top-performing students have more of these skills, they will benefit more from classroom PC use.

Column 1 of Table 4.4 shows the same estimation as Column 3 of Table 4.3, but for a subsample of students whose average test scores in math and science are in the top quartile of the distribution. From an econometric point of view, it is not very compelling to constrain the sample based on the dependent variable. To alleviate this problem, we use the average of the math and science test scores in order to restrict the sample to top-performing students. The most striking result in the top-performance sample is that the effect on science test scores of looking up ideas almost doubles compared to the effect in the full sample. A possible explanation for this finding is that the strongest students are best able to benefit from researching a topic without the teacher's support. As in the full sample, looking up ideas in math has no effect on student achievement. For practicing skills, the negative effect on science remains almost unchanged, while it is almost halved in math. It thus could be concluded that for strong students there is not much harm in practicing math skills on a PC. An observation that is more difficult to explain is that for top-performing students, processing data in science has a negative effect on test scores.

In Column 2 of Table 4.4 we look at students whose average test scores in math and science are in the lowest quartile. While the effects of processing data are very similar for top- and low-performing students, the effects of practicing skills are not. For the top performers, practicing skills in science has a negative effect and but practicing skills in math does not; for the low performers, practicing skills in math has a negative effect, although it is significant only at the 10 percent level. The most interesting result is that the weakest students do not benefit from looking up ideas on a PC; the effect of looking up ideas in science is drastically reduced and the effect on math is actually negative. This suggests that for low-performing students, looking up ideas is not a very effective learning method, especially not in math. The combined effect underlines that due to the negative impacts of two PC activities, PC use in math has a significant negative effect on the test scores of the weakest students. To summarize, it is the strong students who benefit most from looking up ideas in science, whereas for weak students, not a single PC activity has a positive effect on test scores.

Whether classroom use of PCs is beneficial or detrimental for disadvantaged students is a topic of heated debate. Notable examples of empirical studies that focus on the PC use effect for disadvantaged students include Leuven and Lindahl (2007) and Banerjee et al. (2007). Leuven and Lindahl use a regression discontinuity design to analyze the effect of extra computer equipment in Dutch schools in which at least 70 percent of the students have parents with low levels of education. The authors find that the additional computer equipment has no significant effect on boys, but a significant negative effect on girls. Banerjee et al. conduct an experiment in Indian schools that cater to children from especially poor families. They find that two hours of weekly computer use in school, of which one hour was additional to regular classes, had equally positive and significant effects on the achievement of both boys and girls.

In Columns 3 and 4 of Table 4.4, we focus on disadvantaged students along two dimensions: first, we look at students with low socioeconomic status, approximated by the number of books at home, then we turn to non-native students. Column 3 shows results for a subsample of students who belong to the lowest quartile in the distribution of books at home within their country. Compared to the estimation with the full sample of students, there are surprisingly few differences. Both the effects for processing data and the effects for looking up ideas are of a magnitude similar to that found for the full sample. For practicing skills, the negative effects are about 25 percent smaller than in the full sample and not significant. Column 4 shows the results for students who are non-native. Similar to the findings for low-performing students, non-native students do not benefit from using a PC to look up ideas in science, and for math, looking up ideas even has sizable negative effects on



test scores. The combined effect of PC use is negative for both groups of disadvantaged students, although the effect is larger for non-native students.

Another potentially disadvantaged group of students are those who rarely use PCs. It could be that students who frequently use a PC at home benefit more from using a PC at school than do students who are generally less familiar with PCs. There is a large literature on the effects of PC use at home. Most studies, however, find negative or null effects on students achievement. For example, Malamud and Pop-Eleches (2011) analyze the effects of a government program that gave computer vouchers to low-income students in Romania. The vouchers led to a significant negative effect on school grades. Fairlie and Robinson (2013) conduct a randomized experiment in which school children in California were given free computers. Although PC use at home increased substantially in the treatment group, Fairlie and Robinson find no effect on educational outcomes like school grades and student test scores. It has been shown that students with PCs at home do use the computer for schoolwork, but also for gaming, social networking, and many other forms of entertainment. Most studies argue that the potentially beneficial use of PCs for schoolwork is outweighed by less productive PC activity. However, to our knowledge, there has been no study that analyzes whether students who use a PC at home have an advantage in PC activities at school.

In TIMSS, students are asked how often they use a PC at home, with an answer range from "never or almost never" to "every or almost every day." Column 5 of Table 4.4 presents the results for students from the lowest quartile of the PC use distribution within each country. There is surprisingly little difference between the estimates in the low PC use sample and the full sample. In fact, students who rarely use a PC at home benefit to the same degree from looking up ideas in science as does the average student. This is a reassuring result as it appears to mean that students with home computers do not have an advantage over students who do not when it comes to school achievement. There is neither a direct positive effect, as other studies have shown, nor an indirect positive effect due to an advantage in PC activities performed in the classroom.

## 4.6 Country Heterogeneities

Since TIMSS is a cross-country study and our sample includes more than 30 economies, we are able to analyze where PCs are used effectively in schools and where they are not. There is a large literature on the institutional settings that are conducive to good educational outcomes. Hanushek and Woessmann (2011) provide an extensive overview of cross-country studies using

student micro data. One of the earliest studies to systematically investigate the role of institutional differences across countries for student achievement is Woessmann (2003). Woessmann identifies centralized exams, school autonomy, and competition from private schools as factors that are beneficial for student achievement. However, not all institutions have the same effect in every country. For example, Hanushek, Link and Woessmann (2012) show that school autonomy has a positive effect on student achievement in developed countries but a negative effect in developing countries. Similarly, the use of PCs could be conducive to student achievement in some countries and detrimental in other countries.

Maybe the most obvious question regarding country heterogeneities is whether classroom PC use in richer countries is more effective than in poor countries. We are able to answer this question because our sample includes some of the richest countries in the world, such as Qatar, Singapore, Norway, and the United States, as well as relatively poor countries, such as Jordan, Indonesia, Honduras, and Ghana. Previous studies do not provide conclusive evidence as to whether PC use is effective in developing countries. For example, Banerjee et al. (2007) find positive effects in India, but Cristia et al. (2012) find no effects in Peru. There are some theoretical reasons for why PC use could be more effective in richer countries. For one, schools in developed countries have more money to buy appropriate computer equipment and infrastructure. Despite a number of initiatives aimed at providing schoolchildren in developing countries with computers (e.g., the well-known One Laptop per Child Program), schools in these countries might lack basic infrastructure such as electricity, broadband Internet, Wi-Fi routers, and projectors. Also, teachers in developing countries might be less well-trained in how to use PCs in their classes. As our estimates in Tables 4.3 and 4.4 are for the whole sample of countries, we might find stronger effects when restricting the sample to high-income countries.

Column 1 of Table 4.5 presents results for a subsample of countries in the top quartile of the GDP per capita distribution of all TIMSS countries. Contrary to what one might expect, the effect of PC use on student achievement is not generally more positive in rich countries. While using the PC for practicing skills now not have the negative effects we observed in the full sample, neither does looking up ideas in science have the positive effect we observed in the full sample. In fact, the effects of practicing skills and looking up ideas are very close to zero in rich countries. Moreover, processing data in science has a negative effect on test scores, which is why the combined PC use coefficient turns negative in science.

For PCs to be used efficiently in classrooms, teachers need to have appropriate software and digital teaching materials. Educators and policymakers agree that computer hardware alone does no good in the classroom. The applications and materials used for learning at the PC are at least as important as the PC itself. The most challenging task of integrating PCs into instruction is the development of digital content that matches the curriculum. Software developers and publishers have more incentive to design and market applications that complement the curriculum in larger countries. Larger countries also might have a larger community of open-source developers as well as teachers willing to develop free digital content for students. In Column 2 of Table 4.5 we therefore restrict our sample to the countries with the biggest populations. The results are not very different from the estimates for the full sample. The negative effect of practicing skills in science becomes much smaller and the positive effect of looking up ideas in science becomes somewhat smaller but, in general, there is no evidence that countries with a larger population more effectively use PCs in schools.

However, perhaps it is not size of population that matters, but the number of people who speak a country's primary language. Singapore, for example, has a very small population, but students are taught in English. As English is one of the most widely spoken languages in the world, it is safe to assume that there is a great deal of digital school materials that could be used for teaching. Column 3 of Table 4.5 shows results for countries with the largest number of speakers of their primary language. Most estimates are smaller than the estimates in the full sample. There is no evidence that countries with a widely spoken primary language are able to more effectively use PCs in the classroom. Indeed, the results in Columns 2 and 3 show that neither the size of a country nor its number of primary language speakers are decisive for effective PC use.

Today, classroom PCs are usually connected to the Internet. Since many students will need to access the Internet at the same time, a high-speed Internet connection is crucial. To date, only a few studies analyze the effect on student achievement of Internet availability in schools. A notable exception is Belo, Ferreira and Telang (2013), who use a fixed effects and an instrumental variable approach to measure the effect of broadband Internet in Portuguese schools. They find sizable negative effects of broadband on grades in national exams. However, the negative effects are less severe when schools employ strategies to prevent distraction. Particularly, schools that block YouTube significantly dampen the negative effects of Internet availability. Thus, there are two possible roles imaginable that the Internet could play in schools. Faster Internet could either make PC use in the classroom more effective because students do not waste time waiting for websites and online applications to load, or faster Internet could lead to more wasting of time (and thus less

effective PC use) due to the distraction of entertainment websites that can be visited due to the high bandwidth.

Although our data let us know whether schools have Internet, we have no information on the speed of the Internet. Possible proxies for Internet speed in schools are the penetration of broadband Internet and the average Internet speed in a country. Of course, these variables could also capture some other variation that might be relevant for the effectiveness of classroom PC use. For example, a higher broadband penetration could also mean that the country's population has more experience with the Internet. These other channels are equally interesting in our context. Column 4 of Table 4.5 shows estimates for the countries with the highest broadband penetration. There are no large differences between these and the estimates for the full sample. Although the positive effect of PC use for looking up ideas in science is slightly smaller than in the full sample, the effect of looking up ideas in math increases and becomes significant at the 10 percent level. The small differences in the estimates could either mean that broadband penetration does not matter for PC use effectiveness or it could mean that there are competing effects, for example, a positive effect due to more efficient web browsing and a negative effect due to more distraction, which cancel each other out. The results for countries with high average Internet speed are not shown in Table 4.5 due to space limitations, but they are similar to the results for countries with high broadband penetration.

In Column 4, we took a look at countries with a high number of broadband Internet subscriptions, an aspect that might, at least to some extent, capture the Internet experience of the average inhabitant. However, while high-speed Internet is important for schools, it might be less decisive for the individual in that most online applications, such as email or basic web browsing, do not require broadband. For teachers, it thus might be more important to have any kind of Internet access at home, regardless of speed, but to use it frequently. Also, students might be better positioned to use the Internet at school if they have it at home. Therefore, in Column 5 of Table 4.5 we repeat our estimations with a subsample of countries that have a high percentage of Internet users in their population. We observe no major differences in these results compared to the estimates for our full sample. Specifically, there is no indication that classroom PC use is more effective in countries with a high percentage of Internet users.

Overall, the country heterogeneities presented in Table 4.5 do not reveal any clear pattern that answers the question of which kind of countries most effectively use PCs in the classroom. We also explored whether there is any meaningful interaction with several other country characteristics, including educational spending, the number of computers per 100 inhabitants, age structure, average

education and skill level in the population, mobile broadband penetration, and the average use of PCs in schools as measured in the PISA study from 2000. We were not able to identify a single decisive factor that is determinate of effective classroom PC use. In Table 4.5, we look at countries from a specific part of the distribution. To ensure that we have not overlooked a relationship between the effectiveness of PC use and some country characteristic, we estimate the effect of combined PC use in each country and plot the estimates against selected country variables.

Figure 4.1 shows the country coefficients for the combined PC use effect in math and science on the x-axis; the y-axis depicts GDP per capita in the upper panel and total spending on education as a percentage of GDP in the lower panel. None of the four graphs reveal a strong correlation. There could be a small positive correlation with GDP per capita, but educational spending does not seem to be related to combined PC use effectiveness. What the graphs do show is that most countries have small but positive coefficients of combined PC use in science. In math, there are more countries with negative coefficients, most notably Japan and Malaysia with estimates larger than minus 20 percent.

Figure 4.2 presents correlations between PC use effectiveness and computer as well as broadband Internet penetration in the population. There is no evidence of a correlation between the combined PC use coefficients for math and the penetration variables. However, there might be a slightly positive correlation for the combined PC use coefficients in science. As in the previous figure, the science coefficients are much more densely distributed than the math coefficients. This makes it somewhat difficult to spot a correlation, but it is still clear that no strong correlations are present.

Since we find no strong country heterogeneities with respect to the effectiveness of PC use in classrooms, the question arises as to whether there are systematic differences across countries or country characteristics in the way that teachers and students use PCs. For example, do students in richer countries use PCs more often? Is there a relationship between educational spending and the frequency of PC use? If there are no systematic heterogeneities with regard to the intensity of PC use in classrooms, this might to some extent explain why there are no systematic heterogeneities in the effectiveness of PC use.

Figure A4.1 correlates the country share of students who use PCs in math and science at least once per month with GDP per capita and educational spending. In the upper panel of Figure A4.1, we observe a positive correlation between the frequency of PC use and GDP per capita. There are a few outliers, such as Japan, which has a high GDP per capita but relatively little PC use, and Jordan,

which has a very low GDP per capita but high PC use. In general, however, the positive correlation between the two variables is undeniable. The next obvious question is where this positive correlation comes from. Is it simply the greater abundance of resources or is it a country's relatively stronger emphasis on education? The lower two graphs of Figure A4.1 show that the relationship between PC use and relative educational spending is much weaker than the relationship with GDP per capita. Although the trend line still has a positive slope, at least in math it appears that the correlation is mainly driven by Norway and Australia, which have both high PC use and high educational spending.

In Figure A4.2, we show correlations between PC use and the percentage of households with computers as well as the number of broadband Internet subscriptions per 100 inhabitants. The upper two graphs show that PC use in math and science is positively correlated with the percentage of households with computers. Like all measures of ICT infrastructure, the number of computers in households is not independent from GDP per capita. Nevertheless, there is some clear variation between GDP and computer penetration. Malaysia, for example, has a GDP per capita that is half that of Italy's, but computer penetration in the two countries is very similar. The lower two graphs show that there is also a positive correlation between PC use and the number of broadband subscriptions per 100 inhabitants. Again, Japan is a notable outlier, with a high number of broadband subscriptions but only about 20 percent of students using PCs in the classroom at least once per month.

To summarize, we find that the frequency of PC use in the classroom is positively correlated with observable country characteristics such as GDP per capita and computer and broadband penetration. Relative educational spending has a less obvious positive relationship with PC use. Aside from these differences in usage, the effectiveness of PC use is difficult to predict based on country heterogeneities. We therefore speculate that the effectiveness is driven by factors that are difficult to observe and that are not necessarily related to obvious country characteristics. From the education economics literature we know how important teachers' role is in the classroom, which, unfortunately, can be only weakly attributed to observable teacher characteristics (see, e.g., Rivkin, Hanushek and Kain 2005). It is likely, that also the effectiveness of PC use is strongly dependent on the teacher. Factors like ICT infrastructure might therefore play only a subordinated role.

## 4.7 The Fourth-Grade Sample

In the previous sections, we focused on the sample of eighth-grade students. TIMSS has the advantage that it offers data on students in both the eighth grade and the fourth grade. Based on theoretical considerations and previous literature, we assume that older students are able to benefit more from PC use in the classroom, but this does not mean that teachers of lower grades do not use PCs for teaching. In fact, 4.6 shows that there are PCs available in 39 percent of all fourth-grade math classrooms and in 43 percent of all fourth-grade science classrooms. Unfortunately, TIMSS does not provide information for fourth-graders on all three PC activities that are available for eighth-grade students. Processing data is probably too advanced or too specific a task for fourth-graders, leaving us with the activities of practicing skills and looking up ideas. Table 4.6 shows that 26 and 30 percent of students use the PC for practicing skills in science and math, respectively.; 35 and 23 percent of students use the PC for looking up ideas in science and math. As in the eighth-grade sample, there is sufficient variation in PC use on the student level, as shown in Table A4.2.

A subsample of fourth-grade students have the same teacher for both science and math, something that occurs rarely for eighth-graders. By restricting our sample to students who have the same teacher in both subjects, we can use within-student within-teacher variation to identify the effect of PC use on achievement. A potential omitted variable bias due to unobserved teacher characteristics is thereby eliminated. The only remaining concern is that teachers select themselves into classrooms based on the relative performance of students in math and science. For example, teachers who are better at using PCs for math instruction than for science instruction might select themselves into classrooms of students who are better in math than in science. However, this sort of selection does not seem very likely.

Table 4.7 shows results for our baseline specification in the full sample and in the same-teacher sample. The coefficients of combined PC use in the full sample are basically zero and do not change much when controls are added. In the same-teacher sample, there is a positive and statistically significant effect of combined PC use in science. After including controls, this effect becomes drastically smaller and is very close to zero for both subjects. Practicing skills has a small negative effect on math test scores in the full sample. In the same-teacher sample, it also has a negative effect in math, but a small positive effect in science. With controls, however, the positive effect is reduced and only the negative effect in math remains statistically significant. For looking up ideas, the coefficients do not differ much between the two samples and the different specifications. There

is a small but statistically significant positive effect of looking up ideas in math and no effect in science.

Interestingly, the coefficients in the same-teacher sample are more sensitive to adding teaching method controls than are the coefficients in the full sample. Although the differences in coefficients for the single PC activities are not large, they are large enough to change the combined PC use coefficients from a positive effect in science to a null effect. Teachers who teach the same class in math and science are apparently using different teaching methods for each subject that are correlated to PC use. The coefficients in the same-teacher sample could signify that the teacher's use of PCs in science is correlated with the use of other teaching methods that are beneficial for student achievement. This would explain the positive coefficient of PC use in science, which becomes much smaller after controlling for teaching methods. Whether we rely on the results with or without teaching method controls depends on the question we are trying to answer. If the question concerns the effect of PC use as it is usually employed by teachers, the estimates without teaching method controls are preferable. If the question, instead, concerns the isolated effect of PC use, one should probably use the specification that includes all controls. For the remainder of this paper, we focus on estimations in the fourth-grade same-teacher sample with full controls so as to identify the isolated effect of PC use.

Comparing the fourth-grade results to those for the eighth grade reveals many similarities but also a few differences. Generally, practicing skills has negative effects, whereas looking up ideas has positive effects. The differences in the effects between subjects are not the same for eighth-grade and fourth-grade students. In the eighth grade, practicing skills has negative effects in both math and science, while the negative effect in the fourth-grade sample is restricted to math. There is a positive effect of looking up ideas on science for eighth-grade students, but it is in looking up ideas for math that the positive effect is found for fourth-grade students. Regarding the size of the coefficients, effects are generally much smaller in the fourth grade. The negative effects for practicing skills are larger than minus 2 percent of a standard deviation in the eighth-grade sample; the negative effect in math is only half as large in the fourth-grade sample. The positive effect of looking up ideas in science is larger than 3 percent in the eighth-grade sample, while the positive effect of looking up ideas in math is only between 1 and 1.5 percent in the fourth-grade sample.

Table 4.8 presents student heterogeneities in the fourth-grade sample. Column 1 shows that students in the top quartile of combined performance in math and science are not able to benefit drastically more from PC use in the classroom than is the average student. Although the positive



effect for looking up ideas in math increases in the top-performance sample, it is still below 3 percent. Students from the bottom performance quartile in each country do not benefit from using the PC to look up ideas. However, they also do not experience the negative effect of practicing skills in math. For combined PC use, there is a small positive effect in math for low-performing students. This is in contrast with the results for the eighth grade, where top-performing students also benefitted from looking up ideas, but low-performing students were strongly impaired by PC use.

Column 3 of Table 4.8 shows that students with few books at home are, similar to low-performing students, not able to benefit from looking up ideas, but are also less affected by the detrimental effects of practicing skills. Fourth-grade students are not asked their nationality, so we cannot estimate the effects for non-native students as we did for eighth-graders. We do know, however, whether students mostly speak the language of the test at home and we restrict our sample to these students in Model 4. For non-native-language students, there are sizable negative effects of looking up ideas in science, but practicing skills in science has positive effects for these students. Finally, in the case of students who rarely use PCs at home, there are no positive effects from looking up ideas or negative effects from practicing skills.

Fourth-grade disadvantaged students appear to be more negatively affected by classroom PC use than disadvantaged eighth-graders. Students with parents who attained relatively less education, approximated by the books at home variable, do not benefit from using the PC to look up ideas. The same is true for students with low PC use at home. Students who speak a different language at home are severely negatively affected by looking up ideas. In the eighth-grade sample, we did not find these strong heterogeneities. Specifically, students with few books at home and students who rarely use PCs benefit from looking up ideas to a similar degree as does the average student. Thus, it could be concluded that a student's socioeconomic status is more important in regard to effect PC use in the fourth grade than in the eighth grade. On the positive side, disadvantaged students in the fourth-grade sample were not as negatively affected by using the PC to practice skills.

Table 4.9 shows country heterogeneities in the fourth-grade sample. Surprisingly, combined PC use has positive effects in countries that have high GDP per capita, a large number of primary language speakers, and a high rate of broadband subscriptions per 100 inhabitants. These positive combined PC use effects do not come from large coefficients for looking up ideas, as one might expect. Instead, they are driven by positive coefficients for practicing skills, especially practicing skills in science. These results are in contrast to the country heterogeneities in the eighth-grade

sample. Countries with high GDP per capita, a large number of primary language speakers, and high broadband penetration do not have more effective classroom use of PCs for eighth-grade students, but apparently they do use PCs more effectively for fourth-grade students.

## 4.8 Conclusion

This study is the first to systematically analyze the effects of different PC use activities in classrooms. We find small negative and null effects for combined PC use; however, but there are some uses that improve student achievement, namely, using a PC to look up ideas, as well as uses that have a detrimental effect on student achievement, namely, using a PC to practice skills. The opposite effects of different PC use activities offset each other and we speculate that this is a major reason for the inconclusive results of previous studies.

Although PC use is slightly more beneficial for top-performing students and more detrimental for low-performing and non-native students, effect heterogeneities among different groups of students are not very large. We find no evidence that in countries with higher GDP per capita, higher relative educational spending, a larger population, or better ICT infrastructure, PCs are used more effectively by eighth-grade students. For students in the fourth grade, we find stronger heterogeneities with respect to student and country characteristics.

The reason we need to think critically about different PC use activities is not so much because of their potential to negatively impact test scores, but more because of the opportunity cost of time use in the classroom. Using the PC for an activity that does not have positive effects on student learning takes away time from activities that are potentially more effective in improving achievement. This is true for the use of PCs compared to completely different teaching methods, like traditional lecture-style teaching, but also for more and less effective activities on the PC. Although processing and analyzing data does not have negative effects on science test scores, it is still advisable to avoid this activity and instead use the PC for looking up ideas and information.

To assess the effectiveness of an educational innovation, it is not sufficient to look at the potential positive effects and ignore the costs. Even if we take the most positive effects that we find, we need to keep in mind the massive costs of bringing PCs into every classroom. In the United States, for example, there are about 50.1 million students enrolled in public elementary and secondary schools.<sup>6</sup> School districts that recently decided to equip every student and teacher with a computer paid

---

<sup>6</sup> As per the fall of 2013. Source: <http://nces.ed.gov/fastfacts/display.asp?id=372>

between \$344 for Chromebook laptops and \$768 for iPads including new software that is adapted to the curriculum.<sup>7</sup> On top of these costs, school districts have to upgrade their Internet and Wi-Fi infrastructure at schools, which can be more expensive than the new devices themselves.<sup>8</sup>

Let us assume that we want to equip every student and all 3.3 million teachers in public US schools with Chromebook laptops at a cost of \$344 per device. This alone would cost \$18.4 billion. Due to the rapid obsolescence of computer technology, this kind of investment would probably have to be made every three to four years. Spread over 3.5 years, the costs would amount to \$5.3 billion per year, which is certainly a lower bound as many expenditures, like teacher training and salaries for technical staff, not to mention maintenance costs, have not been taken into account. On top of that are infrastructure costs, which can be conservatively estimated to be 50 percent of the device costs, adding another \$9.2 billion to the total costs.

It is difficult to put these costs into perspective with the costs of other teaching methods because many teaching methods have no direct costs at all. In fact, there is a great deal of evidence suggesting that the relationship between student achievement and school resources is very weak (Hanushek, 1997) and that the most effective way to raise student test scores is to raise teacher quality by creating appropriate incentive structures (Hanushek, 2003). Given that most teaching methods can be changed at low cost and that the most effective changes to the education system are not necessarily input based, we compare the cost of PCs in classrooms to other costly measures, independently from their effectiveness.

A typical benchmark for expensive educational reforms is the reduction of class size. The average class size in U.S. public secondary schools was 23.4 pupils in 2008.<sup>9</sup> Reducing class size to 20 pupils would require  $3.4/20 = 17\%$  more classes. The yearly current expenditures of public elementary and secondary schools amounted to about 552.5 billion in 2011.<sup>10</sup> Assuming that the costs of classrooms are proportional to current expenditures (Krueger, 2003), 17 percent more classes would increase current expenditure by 17 percent or \$93.9 billion. Compared to this tremendous cost for a policy measure with very ambiguous effects on student achievement, the cost of equipping every teacher and student with a computer look relatively moderate.

---

<sup>7</sup> <http://www.latimes.com/local/lanow/la-me-ln-laUSD-surveys-ipad-costs-20140112,0,3048105.story>

<sup>8</sup> For example, the Los Angeles Unified School District will pay approximately \$500 million for iPads and software and \$800 million for upgrading Wi-Fi infrastructure. Source: <http://www.scpr.org/blogs/education/2014/02/11/15811/la-schools-wifi-networks-to-cost-about-800-million>

<sup>9</sup> <https://nces.ed.gov/fastfacts/display.asp?id=28>

<sup>10</sup> [http://nces.ed.gov/programs/coe/indicator\\_cmb.asp](http://nces.ed.gov/programs/coe/indicator_cmb.asp)

However, even though equipping every student with a PC is considerably less expensive than reducing class size by about three students, the effects of PC use on student achievement are not exceptionally large. If PC use was restricted to looking up ideas in science, which we found to be the most effective PC activity, switching from not having any PCs to using a PC almost every day would correspond to an improvement in test scores of about 9 percent of a standard deviation. According to Schwerdt and Wuppermann (2011), a similar effect can be achieved by changing the lecture style or, according to Aslam and Kingdon (2011), by quizzing students 30 minutes more per week.

PC use in schools has changed a great deal since the first schools started experimenting with using PCs for instruction in the early 1980s (Oppenheimer, 2003). At the beginning, "digital literacy"-knowing how to handle a PC-was the primary focus of PC use in schools. Later, schools started to set up computer labs that students could use to work with educational software. Today, there is a wide consensus among policymakers that PCs belong in each and every classroom, not just in a lab to be shared by all. The empirical literature followed the evolution of PC use in schools but never actually investigated whether the presence or general usage of PCs in the classroom affects student achievement. Some of the most important questions in this regard can only start to be answered when every student has a PC. For example, what should teachers and students do with the PC? How much time should they spend on single PC use activities? Answering these types of questions will be of paramount value for policymakers and educators.

## Figures and Tables

Figure 4.1

Effect sizes of combined PC use (to process data, practice skills & look up ideas) correlated with GDP per capita and total spending on education




Figure 4.2  
 Effect sizes of combined PC use (to process data, practice skills & look up ideas) correlated with computer and broadband access




Table 4.1  
PC availability and usage intensities by subject for eighth-grade students in TIMSS

	Min.	Max.	Math			Science		
			Mean	SD	Obs.	Mean	SD	Obs.
PCs are available in classroom	0	1	0.342	0.474	183,112	0.427	0.495	180,043
Any PC use for analyzing data	0	1	0.155	0.362	181,339	0.26	0.438	177,991
Any PC use for practicing skills	0	1	0.177	0.382	181,187	0.25	0.433	178,019
Any PC use for looking up ideas	0	1	0.17	0.376	181,327	0.328	0.469	178,263
Analyzing data intensity if PCs are available	1	4	1.662	0.828	59,566	1.887	0.853	72,629
Practicing skills intensity if PCs are available	1	4	1.813	0.922	59,414	1.881	0.878	72,657
Looking up ideas intensity if PCs are available	1	4	1.725	0.84	59,554	2.162	0.84	72,901
Analyzing data intensity incl. classrooms without PCs	0	4	0.554	0.918	181,339	0.791	1.083	177,991
Practicing skills intensity incl. classrooms without PCs	0	4	0.604	1.007	181,187	0.79	1.089	178,019
Looking up ideas intensity incl. classrooms without PCs	0	4	0.575	0.946	181,327	0.908	1.198	178,263

Notes: The table shows the number of observations, the mean, the standard deviation and the minimum and maximum of different PC use related variables in the eighth-grade sample of TIMSS 2011.

Table 4.2  
Cross-sectional OLS and seemingly unrelated regressions for eighth-grade students

	OLS						SUR			
	(1)		(2)		(3)		(4)			
	math	science	math	science	math	science	math	science	math	science
Combined PC use	0.1076*** (0.0194)	0.1812*** (0.0199)	-0.0792*** (0.0210)	-0.0264 (0.0253)	-0.0203 (0.0189)	0.0254 (0.0229)	-0.0294*** (0.0099)	0.0074 (0.0112)		
Processing data	0.0898 (0.0784)	0.0927** (0.0419)	0.0318 (0.0335)	0.0219 (0.0236)	0.0610** (0.0261)	0.0419* (0.0224)	0.0193 (0.0148)	0.0068 (0.0106)		
Practicing skills	0.0891* (0.0497)	0.0062 (0.0402)	-0.1039*** (0.0269)	-0.0026 (0.0247)	-0.0764*** (0.0257)	0.0040 (0.0223)	-0.0390*** (0.0135)	-0.0171 (0.0111)		
Looking up ideas	-0.0682 (0.0533)	0.0879 (0.0535)	-0.0091 (0.0267)	-0.0520* (0.0315)	-0.0060 (0.0260)	-0.0233 (0.0267)	-0.0109 (0.0136)	0.0202 (0.0138)		
Basic controls	-	-	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Teacher controls	-	-	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Teaching controls	-	-	-	-	Yes	Yes	Yes	Yes	Yes	Yes
observations	180,965	177,469	172,099	168,701	168,446	164,832			153,801	
clusters	7,067	6,920	6,661	6,526	6,530	6,386			5,930	

*Notes:* Dependent variables are student test scores in math and science, respectively. The table shows separate estimations in the upper and lower panels. Regressions in the upper panel include an index which measures combined PC use for all PC activities and in the lower panel, the three PC activities enter the regressions simultaneously. The first 6 columns are estimated separately by OLS, the last two columns are estimated by seemingly unrelated regressions (SUR). Standard errors clustered on the classroom level in parentheses. In the SUR model, standard errors are estimated by maximum likelihood. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%



Table 4.3  
Within-student between-subject effect of PC use activities for eighth-grade students

	(1)		(2)		(3)	
	math	science	math	science	math	science
<b>Combined PC use</b>						
Implied $\beta$	0.0046 [0.9433]	-0.0012 [0.0503]	-0.0233** [5.8209]	-0.0019 [0.0282]	-0.0292*** [9.7631]	-0.002 [0.0311]
$\beta_{math} - \beta_{science}$	0.0058		-0.0215		-0.0272**	
$\eta_{math} - \eta_{science}$	0.015		0.0063		0.0001	
<b>Processing data</b>						
Implied $\beta$	0.0103 [0.6471]	-0.007 [0.366]	-0.0137 [0.8092]	-0.0034 [0.1017]	-0.0032 [0.0556]	-0.0082 [0.5354]
$\beta_{math} - \beta_{science}$	0.0172		-0.0102		0.005	
$\eta_{math} - \eta_{science}$	-0.01		-0.0308		-0.0211	
<b>Practicing skills</b>						
Implied $\beta$	-0.034** [6.501]	-0.0317** [6.5457]	-0.0261** [4.0884]	-0.0262** [5.9637]	-0.0206* [2.9774]	-0.0222** [4.1162]
$\beta_{math} - \beta_{science}$	-0.0022		0.0001		0.0016	
$\eta_{math} - \eta_{science}$	0.0847**		0.0909***		0.0695**	
<b>Looking up ideas</b>						
Implied $\beta$	0.0273** [5.0314]	0.0376*** [8.2799]	0.0174 [1.8541]	0.033** [5.8422]	-0.0051 [0.1603]	0.0329** [5.6977]
$\beta_{math} - \beta_{science}$	-0.0103		-0.0156		-0.038**	
$\eta_{math} - \eta_{science}$	-0.0573		-0.058		-0.049	
Basic controls	Yes		Yes		Yes	
Teacher controls	-		Yes		Yes	
Teaching controls	-		-		Yes	
N	157192		153801		153801	
clusters	6053		5930		5930	

Notes: Dependent variables are student test scores in math and science, respectively. Correlated random effects model estimated by seemingly unrelated regressions (SUR) with maximum likelihood and clustered standard errors on the classroom level. The table shows separate estimations in the upper and lower panels, with an index of combined PC use in the upper and three separate PC activity measures as independent variables in the lower panel. Implied  $\beta$  represents the effect of PC use, given by the difference between the estimate of PC use in the respective subject in the equation of the respective subject and the estimate of PC use in the respective subjects in the equation of the other subject (see equations 4 and 5).  $\chi^2$  statistics in brackets. For each variable, the difference between the implied  $\beta$  and  $\eta$  coefficients in math and science are shown. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 4.4  
The effect of PC use activities in subsamples of top-performing and of disadvantaged eighth-grade students

	(1)		(2)		(3)		(4)		(5)	
	Top performance		Low performance		Few books		Non-native		Low PC use	
	math	science	math	science	math	science	math	science	math	science
<b>Combined PC use</b>										
Implied $\beta$	-0.0152 [1.1818]	-0.0222 [2.4225]	-0.0438*** [11.8246]	-0.0052 [0.1382]	-0.0199* [3.6092]	0.0032 [0.0681]	-0.0573*** [7.2893]	-0.0065 [0.0571]	-0.0182 [2.646]	0.0107 [0.7761]
$\beta_{math} - \beta_{science}$	0.007		-0.0387**		-0.0231		-0.0508		-0.0289*	
$\eta_{math} - \eta_{science}$	0.069**		-0.0191		-0.0085		0.0395		-0.0127	
<b>Processing data</b>										
Implied $\beta$	-0.004 [0.0485]	-0.0354** [4.8666]	0.0274 [1.7905]	-0.0266* [2.822]	0.0101 [0.4272]	-0.0048 [0.1577]	0.0682** [3.9595]	-0.0109 [0.2384]	-0.0046 [0.1018]	0.0018 [0.0229]
$\beta_{math} - \beta_{science}$	0.0314		0.054**		0.0149		0.0791*		-0.0065	
$\eta_{math} - \eta_{science}$	-0.0278		0.0146		-0.0312		0.1149*		0.0106	
<b>Practicing skills</b>										
Implied $\beta$	-0.0115 [0.4461]	-0.0282** [4.2413]	-0.0333* [3.3167]	0.0027 [0.0351]	-0.0178 [1.5442]	-0.0192 [2.6667]	-0.012 [0.1774]	-0.0284 [1.2827]	-0.0179 [1.7182]	-0.0231* [3.604]
$\beta_{math} - \beta_{science}$	0.0168		-0.036		0.0014		0.0164		0.0052	
$\eta_{math} - \eta_{science}$	0.0942***		0.032		0.0507		0.0937		0.0419	
<b>Looking up ideas</b>										
Implied $\beta$	0.0027 [0.0248]	0.0519*** [7.9986]	-0.0424** [5.1247]	0.0194 [1.1308]	-0.0129 [0.7689]	0.0312** [4.7142]	-0.1159*** [13.9333]	0.0321 [0.8838]	0.0051 [0.1408]	0.0389** [6.2162]
$\beta_{math} - \beta_{science}$	-0.0493**		-0.0618**		-0.0442**		-0.148***		-0.0338*	
$\eta_{math} - \eta_{science}$	0.0013		-0.0673*		-0.0262		-0.1842***		-0.0712*	
Basic controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Teacher controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Teaching controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	37341	37341	36591	36591	63091	63091	20281	20281	70355	70355
clusters	4722	4722	4703	4703	5606	5606	4315	4315	5589	5589

Notes: Dependent variables are student test scores in math and science, respectively. Correlated random effects model estimated by seemingly unrelated regressions (SUR) with maximum likelihood and clustered standard errors on the classroom level. Subsamples (1) and (2) are based on the top and bottom quartile of the distribution of combined test scores in math and science per country. Subsamples (3) and (5) are based on the bottom quartile of the distribution of the respective variable per country. The table shows separate estimations in the upper and lower panels, with an index of combined PC use in the upper and three separate PC activity measures as independent variables in the lower panel. Implied  $\beta$  represents the effect of PC use, given by the difference between the estimate of PC use in the respective subject in the equation of the respective subject and the estimate of PC use in the respective subjects in the equation of the other subject (see equations 4 and 5).  $\chi^2$  statistics in brackets. For each variable, the difference between the implied  $\beta$  and  $\eta$  coefficients in math and science are shown. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 4.5  
The effect of PC use activities for eighth-grade students in subsamples of countries with specific characteristics

	(1)		(2)		(3)		(4)		(5)	
	math	science	math	science	math	science	math	science	math	science
<b>Combined PC use</b>										
Implied $\beta$	0.0025 [0.0428]	-0.0248** [4.1923]	-0.0325*** [8.3273]	0.0001 [0.0001]	-0.0127 [1.3866]	-0.0075 [0.444]	0.0111 [0.7544]	-0.0247** [4.2066]	0.0058 [0.1231]	-0.0228 [2.2419]
$\beta_{math} - \beta_{science}$	0.0273	-0.0327**	-0.0052				0.0358**		0.0286	
$\eta_{math} - \eta_{science}$	0.0381	-0.0202	0.0058				0.0627*		0.0892**	
<b>Processing data</b>										
Implied $\beta$	0.0184 [1.0594]	-0.0265** [4.9017]	-0.0114 [0.3912]	-0.0155 [1.2549]	-0.0014 [0.0063]	-0.0124 [0.92]	-0.0124 [0.4285]	-0.0128 [1.0357]	-0.013 [0.371]	-0.0216 [1.6733]
$\beta_{math} - \beta_{science}$	0.0448**	0.0041	0.0109				0.0004		0.0087	
$\eta_{math} - \eta_{science}$	0.0036	-0.0262	0.0414				0.0153		0.0169	
<b>Practicing skills</b>										
Implied $\beta$	-0.0074 [0.2061]	-0.0002 [0.0004]	-0.0212 [1.601]	-0.0036 [0.0743]	-0.0017 [0.011]	-0.0041 [0.1086]	-0.006 [0.1565]	-0.027** [4.3444]	0.0061 [0.1313]	-0.0244 [2.6866]
$\beta_{math} - \beta_{science}$	-0.0072	-0.0175	0.0024				0.021		0.0305	
$\eta_{math} - \eta_{science}$	0.0393	0.0594	0.0055				0.0725*		0.0155	
<b>Looking up ideas</b>										
Implied $\beta$	-0.0111 [0.415]	0.0039 [0.0544]	-0.0005 [0.0008]	0.0219 [1.4965]	-0.0108 [0.4234]	0.0123 [0.6784]	0.0325* [3.1462]	0.0207 [1.7949]	0.0131 [0.3879]	0.0311* [2.7239]
$\beta_{math} - \beta_{science}$	-0.015	-0.0225	-0.0232				0.0118		-0.018	
$\eta_{math} - \eta_{science}$	-0.0145	-0.0577	-0.0509				-0.0379		0.0679	
Basic controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Teacher controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Teaching controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	26248	35986	32376				28297		21240	
clusters	1369	1360	1278				1272		911	

Notes: Dependent variables are student test scores in math and science, respectively. Correlated random effects model estimated by seemingly unrelated regressions (SUR) with maximum likelihood and clustered standard errors on the classroom level. Subsamples (1)-(5) are based on countries that belong to the top quartile of the distribution of the respective variable in the full sample of countries. The table shows separate estimations in the upper and lower panels, with an index of combined PC use in the upper and three separate PC activity measures as independent variables in the lower panel. Implied  $\beta$  represents the effect of PC use, given by the difference between the estimate of PC use in the respective subject in the equation of the respective subject and the estimate of PC use in the respective subjects in the equation of the other subject (see equations 4 and 5).  $\chi^2$  statistics in brackets. For each variable, the difference between the implied  $\beta$  and  $\eta$  coefficients in math and science are shown. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 4.6  
PC availability and usage intensities by subject for the full and the same-teacher sample of fourth-grade students

	Math				Science			
	Min.	Max.	Mean	SD	Obs.	Mean	SD	Obs.
<b>Full sample</b>								
PCs are available in classroom	0	1	0.386	0.487	268,753	0.431	0.495	264,504
Any PC use for practicing skills	0	1	0.296	0.456	265,709	0.262	0.440	262,372
Any PC use for looking up ideas	0	1	0.231	0.422	265,655	0.349	0.477	262,329
Practicing skills intensity if PCs are available	1	4	2.359	0.967	97,181	1.947	0.922	111,613
Looking up ideas intensity if PCs are available	1	4	1.932	0.911	97,127	2.218	0.821	111,570
Practicing skills intensity incl. classrooms without PCs	0	4	0.900	1.292	265,709	0.830	1.136	262,372
Looking up ideas intensity incl. classrooms without PCs	0	4	0.737	1.094	265,655	0.945	1.221	262,329
<b>Same-teacher sample</b>								
PCs are available in classroom	0	1	0.397	0.489	179,900	0.431	0.495	177,557
Any PC use for practicing skills	0	1	0.313	0.464	178,109	0.269	0.444	176,170
Any PC use for looking up ideas	0	1	0.245	0.430	178,052	0.353	0.478	176,162
Practicing skills intensity if PCs are available	1	4	2.414	0.954	70,103	1.974	0.912	76,499
Looking up ideas intensity if PCs are available	1	4	1.963	0.906	70,046	2.238	0.812	76,491
Practicing skills intensity incl. classrooms without PCs	0	4	0.947	1.321	178,109	0.841	1.143	176,170
Looking up ideas intensity incl. classrooms without PCs	0	4	0.771	1.114	178,052	0.953	1.227	176,162

Notes: The table shows the minimum and maximum, the number of observations, the mean and the standard deviation of different PC use related variables in the full fourth-grade sample and in the sample of fourth-grade students that have the same teacher in math and science.

Table 4.7  
Within-student between-subject effect of PC use activities for fourth-grade students

	Full sample			Same-teacher sample				
	(1)	(2)	(3)	(4)	(3)	(4)		
	math	science	math	science	math	science		
<b>Combined</b>								
Implied $\beta$	-0.0018 [0.069]	0.0017 [0.0516]	0.0013 [0.0369]	-0.0057 [0.6471]	0.0032 [0.2022]	0.0218*** [8.1181]	0.0044 [0.4055]	0.0061 [0.6705]
$\beta_{math} - \beta_{science}$	-0.0035	0.007		-0.0186**			-0.0017	
$\eta_{math} - \eta_{science}$	0.0414	0.0606*		-0.0161			0.0171	
<b>Practicing skills</b>								
Implied $\beta$	-0.0132** [3.9832]	0.0049 [0.4986]	-0.0113* [3.2506]	0.0025 [0.1465]	-0.016** [5.8975]	0.0184*** [7.1449]	-0.0126** [3.9091]	0.0101 [2.2574]
$\beta_{math} - \beta_{science}$	-0.0181**	-0.0139		-0.0344***			-0.0227***	
$\eta_{math} - \eta_{science}$	0.0286	0.0316		-0.0144			0.0039	
<b>Looking up ideas</b>								
Implied $\beta$	0.0103* [3.0065]	-0.0031 [0.1212]	0.0116** [3.8825]	-0.0088 [1.1031]	0.0175*** [9.0076]	0.0038 [0.2215]	0.0154*** [6.7948]	-0.0043 [0.3061]
$\beta_{math} - \beta_{science}$	0.0134	0.0204**		0.0137			0.0197**	
$\eta_{math} - \eta_{science}$	0.013	0.0309		0.0012			0.0162	
Basic controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Teacher controls	-	Yes	Yes	-	-	-	-	-
Teaching controls	-	Yes	Yes	-	-	-	Yes	Yes
N	188374	188374	188374	154291	154291	154291	154291	154291
clusters	8742	8742	8742	7537	7537	7537	7537	7537

Notes: Dependent variables are student test scores in math and science, respectively. Correlated random effects model estimated by seemingly unrelated regressions (SUR) with maximum likelihood and clustered standard errors on the classroom level. The table shows separate estimations in the upper and lower panels, with an index of combined PC use in the upper and three separate PC activity measures as independent variables in the lower panel. (1) and (2) are estimations in the full fourth-grade sample and (3) and (4) are estimations in the sample of fourth-grade students that have the same teacher in math and science. Implied  $\beta$  represents the effect of PC use, given by the difference between the estimate of PC use in the respective subject in the equation of the respective subject and the estimate of PC use in the equation of the other subject (see equations 4 and 5).  $\chi^2$  statistics in brackets. For each variable, the difference between the implied  $\beta$  and  $\eta$  coefficients in math and science are shown. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 4.8  
The effect of PC use activities in subsamples of top-performing and of disadvantaged fourth-grade students

	(1)		(2)		(3)		(4)		(5)	
	Top performance		Low performance		Few books		Non-native language		Low PC use	
	math	science	math	science	math	science	math	science	math	science
<b>Combined</b>										
Implied $\beta$	0.0099	0.0184	0.0235**	0.0085	-0.0033	0.0001	0.0113	-0.01	-0.0036	0.0078
	[0.6975]	[2.1581]	[4.7877]	[0.4221]	[0.1454]	[0.0002]	[0.8146]	[0.6179]	[0.1697]	[0.7078]
$\beta_{math} - \beta_{science}$	-0.0086	0.015			-0.0034		0.0212		-0.0114	
$\eta_{math} - \eta_{science}$	0.0294	-0.0254			0.0021		0.0812		0.0056	
<b>Practicing skills</b>										
Implied $\beta$	-0.0183*	0.0181	0.0155	0.015	-0.0061	0.0074	-0.0065	0.0243**	-0.0115	0.0072
	[2.7657]	[2.4676]	[1.8968]	[1.8802]	[0.4667]	[0.6452]	[0.3012]	[4.0081]	[1.7646]	[0.6344]
$\beta_{math} - \beta_{science}$	-0.0363**	0.0005			-0.0135		-0.0308*		-0.0187	
$\eta_{math} - \eta_{science}$	0.0187	-0.021			-0.0024		0.055		-0.001	
<b>Looking up ideas</b>										
Implied $\beta$	0.0258***	0.0009	0.0081	-0.0093	0.0019	-0.0087	0.0151	-0.0409***	0.007	0.001
	[6.7856]	[0.005]	[0.6656]	[0.5509]	[0.0548]	[0.6638]	[2.0569]	[8.5331]	[0.8161]	[0.0098]
$\beta_{math} - \beta_{science}$	0.0248*	0.0175			0.0106		0.0559***		0.006	
$\eta_{math} - \eta_{science}$	0.0107	-0.0075			0.0018		0.0209		0.0081	
Basic controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Teacher controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Teaching controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	39827	37581	64986	38599	67457					
clusters	6528	6500	7346	6286	7417					

Notes: Dependent variables are student test scores in math and science, respectively. Correlated random effects model estimated by seemingly unrelated regressions (SUR) with maximum likelihood and clustered standard errors on the classroom level. All estimations are with fourth-grade students that have the same teacher in math and science. Subsamples (1) and (2) are based on the top and bottom quartile of the distribution of combined test scores in math and science per country. Subsamples (3) and (5) are based on the bottom quartile of the distribution of the respective variable per country. The table shows separate estimations in the upper and lower panels, with an index of combined PC use in the upper and three separate PC activity measures as independent variables in the lower panel. Implied  $\beta$  represents the effect of PC use, given by the difference between the estimate of PC use in the respective subject in the equation of the respective subject and the estimate of PC use in the respective subjects in the equation of the other subject (see equations 4 and 5).  $\chi^2$  statistics in brackets. For each variable, the difference between the implied  $\beta$  and  $\eta$  coefficients in math and science are shown. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 4.9  
The effect of PC use activities for fourth-grade students in subsamples of countries with specific characteristics

	(1)		(2)		(3)		(4)		(5)	
	math	science	math	science	math	science	math	science	math	science
<b>Combined</b>										
Implied $\beta$	0.0219*** [7.3916]	0.0158* [3.4231]	0.0023 [0.0782]	0.0141 [2.6324]	0.0167* [2.9934]	0.0145 [2.0113]	0.0128 [2.2442]	0.0187** [4.0436]	-0.0067 [0.3352]	-0.0206 [2.6675]
$\beta_{math} - \beta_{science}$	0.0061	-0.0118				0.0023		-0.0059	0.0139	
$\eta_{math} - \eta_{science}$	0.0413	0.0144				0.0109		0.0138	0.0689*	
<b>Practicing skills</b>										
Implied $\beta$	0.015** [3.9541]	0.0259*** [10.4651]	-0.008 [1.1973]	0.0197** [6.3631]	0.0184** [5.0925]	0.0242*** [6.8627]	-0.0006 [0.0061]	0.0187** [4.4588]	-0.0206** [3.9964]	-0.0133 [1.3924]
$\beta_{math} - \beta_{science}$	-0.0109	-0.0277***				-0.0058		-0.0193*	-0.0073	
$\eta_{math} - \eta_{science}$	0.0111	-0.0019				-0.0124		0.0069	0.057*	
<b>Looking up ideas</b>										
Implied $\beta$	0.0076 [1.166]	-0.0148 [2.574]	0.0096 [1.8865]	-0.0077 [0.7232]	-0.0006 [0.0058]	-0.0151 [1.9744]	0.0135* [3.5387]	-0.0013 [0.0167]	0.0115 [1.3905]	-0.0065 [0.2385]
$\beta_{math} - \beta_{science}$	0.0224**		0.0172*			0.0144		0.0148	0.0179	
$\eta_{math} - \eta_{science}$	0.0415	0.0224				0.0321		0.011	0.0096	
Basic controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Teacher controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Teaching controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	35029	50452	50452	31215	31215	31215	36313	27676		
clusters	1968	2226	2226	1523	1523	1523	1872	1451		

Notes: Dependent variables are student test scores in math and science, respectively. Correlated random effects model estimated by seemingly unrelated regressions (SUR) with maximum likelihood and clustered standard errors on the classroom level. All estimations are with fourth-grade students that have the same teacher in math and science. Subsamples (1)-(5) are based on countries that belong to the top quartile of the distribution of the respective variable in the full sample of countries. The table shows separate estimations in the upper and lower panels, with an index of combined PC use in the upper and three separate PC activity measures as independent variables in the lower panel. Implied  $\beta$  represents the effect of PC use, given by the difference between the estimate of PC use in the respective subject in the equation of the respective subject and the estimate of PC use in the respective subjects in the equation of the other subject (see equations 4 and 5).  $\chi^2$  statistics in brackets. For each variable, the difference between the implied  $\beta$  and  $\eta$  coefficients in math and science are shown. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

## Appendix

Figure A4.1

Correlations of GDP per capita and total spending on education with the percentage of 8th grade students using e-learning at least once a month

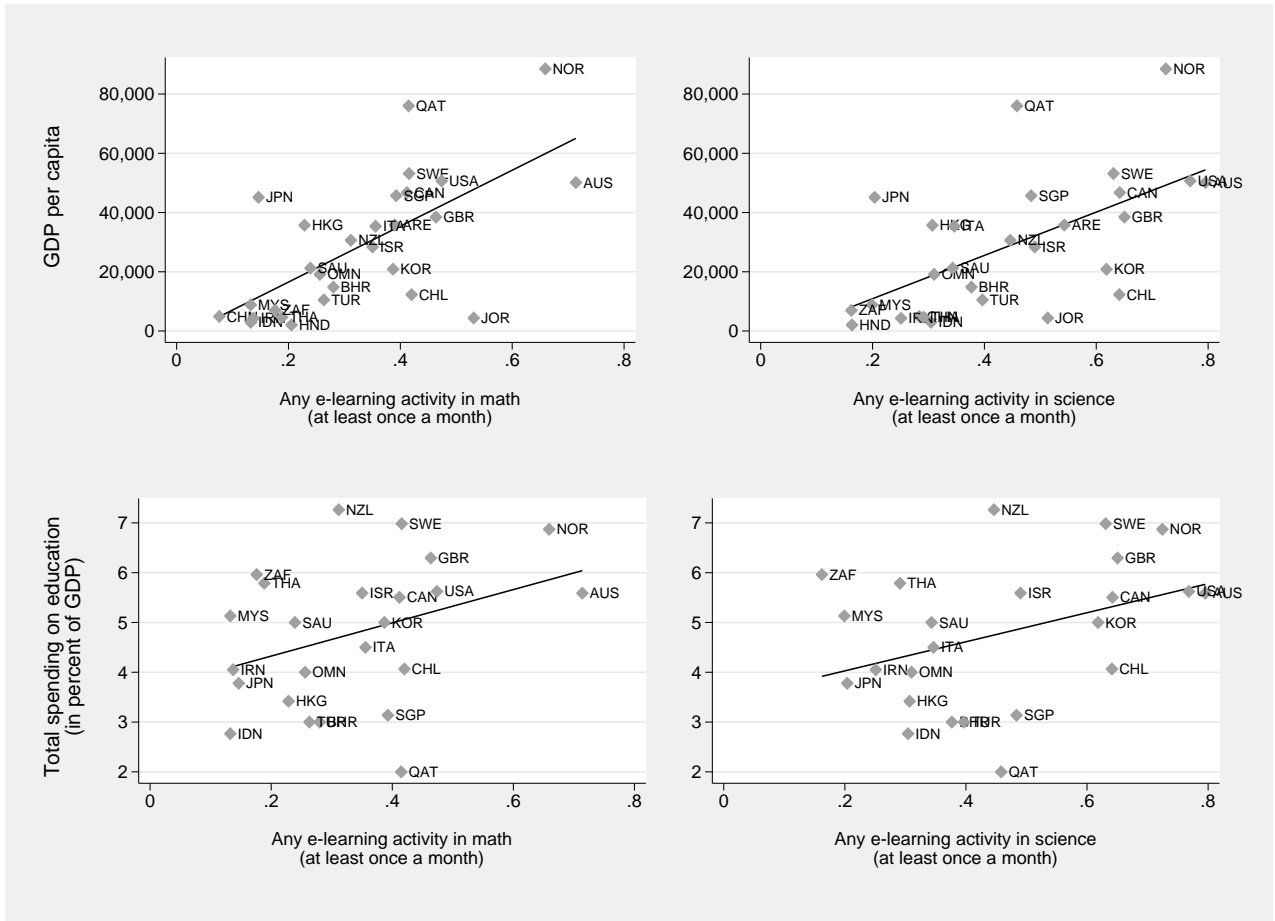




Figure A4.2  
 Correlations of computer and broadband access by country with the percentage of 8th grade students using e-learning at least once a month




Table A4.1  
Share of eighth-grade students with between-subject variation in PC use intensities

Country	Analyzing data			Practicing skills			Looking up ideas		
	Total Obs.	With variation	Variation share	With variation	Variation share	With variation	Variation share	With variation	Variation share
Australia	7,170	5,858	0.82	5,979	0.83	6,068	0.85		
Bahrain	4,640	2,260	0.49	2,237	0.48	2,337	0.50		
Botswana	5,396	1,933	0.36	1,859	0.34	1,892	0.35		
Canada (Alberta)	4,750	3,044	0.64	3,110	0.65	3,277	0.69		
Canada (Ontario)	4,716	2,800	0.59	2,790	0.59	3,018	0.64		
Canada (Quebec)	5,488	3,472	0.63	3,475	0.63	3,599	0.66		
Chile	5,835	4,130	0.71	4,032	0.69	4,283	0.73		
Chinese Taipei	4,995	2,236	0.45	2,312	0.46	2,332	0.47		
England	2,880	1,899	0.66	1,935	0.67	1,924	0.67		
Ghana	6,799	1,466	0.22	1,390	0.20	1,361	0.20		
Honduras	4,320	1,282	0.30	1,374	0.32	1,282	0.30		
Hong Kong	3,881	1,773	0.46	1,803	0.46	1,677	0.43		
Indonesia	2,447	1,055	0.43	1,016	0.42	974	0.40		
Iran	6,029	1,643	0.27	1,798	0.30	1,679	0.28		
Israel	4,263	2,833	0.66	2,874	0.67	2,827	0.66		
Italy	3,979	1,377	0.35	1,361	0.34	1,319	0.33		
Japan	3,955	1,890	0.48	1,828	0.46	2,102	0.53		
Jordan	7,694	4,403	0.57	4,318	0.56	4,496	0.58		
Malaysia	5,733	1,717	0.30	1,694	0.30	1,717	0.30		
New Zealand	5,336	2,784	0.52	2,879	0.54	2,887	0.54		
Norway	3,752	2,027	0.54	1,931	0.51	2,348	0.63		
Oman	9,542	4,198	0.44	4,228	0.44	4,417	0.46		
Palestinian National Authority	7,812	4,433	0.57	4,226	0.54	4,520	0.58		
Qatar	4,422	2,536	0.57	2,562	0.58	2,516	0.57		
Saudi Arabia	4,344	1,764	0.41	1,748	0.40	1,762	0.41		
Singapore	5,927	3,945	0.67	4,119	0.69	3,973	0.67		
South Africa	11,969	4,786	0.40	4,800	0.40	4,747	0.40		
South Korea	4,267	3,242	0.76	3,196	0.75	3,219	0.75		
Sweden	3,155	2,171	0.69	2,187	0.69	2,326	0.74		
Thailand	6,124	2,538	0.41	2,534	0.41	2,677	0.44		
Turkey	6,928	3,426	0.49	3,576	0.52	3,540	0.51		
United Arab Emirates	12,092	7,591	0.63	7,540	0.62	7,389	0.61		
United Arab Emirates (Dubai)	4,847	3,175	0.66	3,104	0.64	2,999	0.62		
United Arab Emirates (Abu Dhabi)	3,727	2,115	0.57	2,192	0.59	2,176	0.58		
United States	10,304	8,602	0.83	8,535	0.83	8,706	0.84		

Notes: The table shows the number of observations for each economy included in the sample in the second column. In the third to eighth column, the table shows the number and the share of observations that have different PC use intensities in math and science for the respective PC use activity.

Table A4.2

Share of fourth-grade students with between-subject variation in PC use intensities

Country	Total observations			Share of observations with variation			
	Full sample	Same-teacher	Share	Full sample		Same-teacher sample	
				Practice	Lookup	Practice	Lookup
Armenia	5146	5146	1.00	0.42	0.43	0.42	0.43
Australia	5741	5043	0.88	0.75	0.71	0.73	0.68
Austria	4610	4366	0.95	0.53	0.51	0.52	0.50
Azerbaijan	4882	3022	0.62	0.40	0.40	0.23	0.24
Bahrain	4083	88	0.02	0.46	0.51	0.38	0.38
Belgium (Flemish)	4849	4849	1.00	0.62	0.68	0.62	0.68
Botswana	4198	2506	0.60	0.27	0.27	0.20	0.20
Canada (Alberta)	3585	3329	0.93	0.62	0.66	0.60	0.65
Canada (Ontario)	4476	3953	0.88	0.48	0.51	0.43	0.46
Canada (Quebec)	4210	2779	0.66	0.52	0.54	0.40	0.43
Chile	5585	5585	1.00	0.62	0.59	0.62	0.59
Chinese Taipei	4284	161	0.04	0.65	0.68	0.54	0.54
Croatia	4584	4584	1.00	0.10	0.12	0.10	0.12
Czech Republic	4578	3558	0.78	0.45	0.41	0.40	0.34
Denmark	3987	1671	0.42	0.73	0.69	0.53	0.50
England	2802	2209	0.79	0.65	0.62	0.61	0.56
Finland	4379	4118	0.94	0.42	0.53	0.40	0.52
Georgia	4799	3422	0.71	0.27	0.26	0.22	0.21
Germany	3779	2064	0.55	0.56	0.54	0.47	0.43
Honduras	3678	3678	1.00	0.17	0.17	0.17	0.17
Hong Kong	3908	469	0.12	0.70	0.71	0.17	0.25
Hungary	5204	4000	0.77	0.32	0.36	0.29	0.31
Iran	5760	5760	1.00	0.07	0.08	0.07	0.08
Ireland	4560	4560	1.00	0.46	0.41	0.46	0.41
Italy	4200	2853	0.68	0.35	0.34	0.31	0.30
Japan	3743	2346	0.63	0.42	0.50	0.27	0.31
Kazakhstan	4382	4382	1.00	0.38	0.32	0.38	0.32
Kuwait	4142	48	0.01	0.44	0.46	1.00	1.00
Lithuania	4688	4577	0.98	0.28	0.31	0.28	0.31
Malta	1703	1038	0.61	0.52	0.48	0.39	0.32
Morocco	7841	1364	0.17	0.45	0.45	0.43	0.43
Netherlands	2671	2671	1.00	0.86	0.71	0.86	0.71
New Zealand	5561	5126	0.92	0.81	0.67	0.81	0.65
Northern Ireland	3546	3448	0.97	0.64	0.52	0.63	0.51
Norway	3119	1901	0.61	0.66	0.57	0.61	0.48
Oman	10347	4272	0.41	0.28	0.27	0.17	0.17
Poland	5027	5027	1.00	0.11	0.13	0.11	0.13
Portugal	4042	4042	1.00	0.29	0.27	0.29	0.27
Qatar	4117	413	0.10	0.57	0.56	0.32	0.32
Romania	4607	4607	1.00	0.20	0.18	0.20	0.18
Russian Federation	4467	4313	0.97	0.15	0.19	0.14	0.18
Saudi Arabia	4515	52	0.01	0.38	0.36	0.42	0.42
Serbia	4379	4379	1.00	0.11	0.11	0.11	0.11
Singapore	6220	3382	0.54	0.57	0.58	0.43	0.45
Slovak Republic	5604	3697	0.66	0.34	0.36	0.26	0.27
Slovenia	4492	4459	0.99	0.34	0.29	0.33	0.29
South Korea	4245	3899	0.92	0.30	0.30	0.26	0.26
Spain	4183	4183	1.00	0.37	0.39	0.37	0.39
Sweden	4597	3131	0.68	0.74	0.68	0.69	0.60
Thailand	4448	4448	1.00	0.30	0.29	0.30	0.29
Tunisia	4912	3463	0.71	0.17	0.18	0.13	0.15
Turkey	7479	7479	1.00	0.23	0.22	0.23	0.22
United Arab Emirates	13817	2512	0.18	0.52	0.52	0.41	0.39
UAE (Abu Dhabi)	4079	292	0.07	0.49	0.50	0.39	0.33
UAE (Dubai)	5497	1903	0.35	0.56	0.57	0.40	0.41
United States	12489	10199	0.82	0.67	0.58	0.62	0.52
Yemen	8058	1269	0.16	0.33	0.34	0.31	0.32

Notes: The table shows the number of observations for each economy included in the full fourth-grade sample and in the the sample of fourth-grade students that have the same teacher in math and science. It also shows the share of students that have different PC use intensities in math and science for the "looking up ideas and information" and the "practice skills and procedures" variables in both samples.



## Chapter 5

---

# The Effect of Mobile Internet Infrastructure on Local Online Services

---

Policymakers call for more investment in broadband Internet infrastructure. Network providers, however, are reluctant to upgrade their network unless demand justifies the additional investment. For this reason, many regulators subsidize Internet infrastructure investment in areas with low demand. In the United States, for example, the Federal Communications Commission (FCC) created the Connect America Fund, which subsidizes broadband deployment in rural areas with up to \$4.5 billion per year (FCC, 2011). A portion of this fund is specifically targeted at expanding mobile broadband infrastructure. In Germany, the regulator forced winners of the latest spectrum auction to roll out their new network in rural areas first (Bundesnetzagentur, 2010). Despite these regulatory efforts, the empirical evidence on the effect of infrastructure upgrades that are not driven by demand is anything but conclusive. Returns are likely to be small if the new infrastructure is either not used at all or used in the same fashion as the older infrastructure. It is therefore important to determine how an upgrade affects usage.

This paper shows that upgrading the mobile Internet infrastructure in Germany from a second-generation network (2G) to a third-generation network (3G) had a positive effect on the usage of local online services. While mobile Internet is useful for a wide variety of applications, such as online search or social networks, it is indispensable for local online services that offer information based on the user's location. It is thus not surprising that travel and navigation apps are among the

most frequently used smartphone applications.<sup>1</sup> Some of the most popular local online services allow users to find restaurants or other businesses based on ratings and reviews by previous customers. The importance of these reviews for restaurants is underlined by Anderson and Magruder (2012) who show that a better rating on the review platform Yelp leads to a significant increase in bookings.

Combining rich data on restaurants from Europe's largest review platform Qype,<sup>2</sup> a unique dataset of all cellular antennas in Germany, and a fine-grained digital surface model (DSM), I find that in areas where for geographical reasons 3G was available early on, restaurants received more reviews and a larger share of these reviews were written on a smartphone. I exploit a technical particularity that affects how 2G and 3G signals are differently absorbed by the terrain. Since the actual availability of 3G at a given restaurant cannot be observed, I estimate the reduced-form effect of specific geographic variation on reviews in a difference-in-difference setting.

While I find positive effects of 3G availability on the number of mobile and non-mobile reviews, I do not find evidence that reviews submitted from mobile phones are considerably more or less favorable than reviews submitted from PCs. If anything, mobile reviews are associated with a slightly better rating than non-mobile reviews. In order to identify differences in rating between mobile and non-mobile reviews, I use both within-restaurant and within-reviewer variation. Although the overall effects are very small, I find evidence that reviews submitted from BlackBerry smartphones are slightly more favorable than reviews from Android smartphones. Reviews from iPhones are the least favorable among all mobile reviews. Due to the very small differences in rating between mobile and non-mobile reviews, I do not find that restaurants with more mobile reviews have a better or worse average rating. Nevertheless, the increased information content of a larger number of reviews are likely to enhance market transparency and promote competition as well as consumer welfare.

The identification strategy that I use for causal inference on the effect of 3G availability is twofold. I use a difference-in-difference model which compares restaurants with high and low likelihood of 3G reception before and after 3G became popular. In order to avoid differences in group composition and trends of the treatment and control group, I exploit exogenous variation in 3G availability to determine the likelihood of 3G coverage in the surroundings of a restaurant. This variation is based on the relationship between the frequency of cellular signals and their sensitivity to other objects. All signals are to some extent absorbed by the terrain, but the degree of absorption

---

<sup>1</sup> Market research shows that after social networks and search, travel apps are the third most popular app category (The Economist, 2012).

<sup>2</sup> In October 2012, Qype was acquired by its largest competitor Yelp and all Qype reviews were migrated to Yelp by October 2013, when the Qype website shut down.

depends on the frequency of the signal (Klemens, 2010, p. 25).<sup>3</sup> The probability that an object completely absorbs a signal is also dependent on the frequency of the signal. For example, it is much harder to see around objects than to hear around objects. This is because the frequency of light is much higher than the frequency of sound. In a similar way, a 3G signal on a high-frequency band is less likely to travel around objects than does a 2G signal on a lower frequency band.

When the 2G infrastructure was built in the early 1990s, every network provider chose its own antenna locations with the aim of maximizing coverage across the country.<sup>4</sup> In the early 2000s, most of Germany was covered by the 2G network. When the 3G infrastructure was rolled out, network providers began to upgrade their existing antenna locations with new 3G antennas. Of course, the roll-out itself was not random and urban areas had priority over rural areas. Nevertheless, there were many areas with good 2G reception, but even when 3G antennas were installed at their closest antenna mast, they were not covered by the 3G signal due to characteristics of the terrain around them. Other areas with equal distance to the next antenna mast, however, were able to benefit from 3G because there were no hills or forests blocking the signal. I therefore compare restaurants in areas with favorable topographic conditions with restaurants in areas that are unlikely to have 3G reception due to unfavorable topographic conditions. To determine the likelihood of 3G reception, I calculate if there is a clear line-of-sight between the area around a restaurant and an antenna mast. If hills, forests, large buildings, or other obstacles are between a restaurant and the next antenna masts, the likelihood of 3G reception drops considerably.

The location-based service I analyze in this study, Qype, is especially useful for mobile Internet users who are looking for restaurants close to their current location. When a tourist, for example, visits a new city, he can open the Qype app on his smartphone and search for nearby restaurants that have received good reviews from previous customers. If he enjoys the food at a chosen restaurant, he can use the same app to rate the restaurant himself and write a review. Alternatively, he can also write a review from his PC at home. While writing reviews is not necessarily bound to the location of a restaurant, the discovery of new restaurants with the help of the smartphone app clearly is. Since in most cases, smartphone users will use the app to find restaurants nearby, I use the immediate surrounding of a restaurant as the basis for the line-of-sight calculations.

---

<sup>3</sup> This relationship is also exploited by Strömberg (2004) who uses a county's share of woodland as an instrument for the share of households with radios. His underlying reasoning is that the radio signal is highly absorbed by the woodland and thus people in woody areas are less likely to have radio reception. Similarly, Nolen and Klonner (2010) use terrain ruggedness as an instrument for the rollout of a cellular network in South Africa.

<sup>4</sup> In Germany, the largest network provider at that time was the state-owned Deutsche Bundespost, which had universal service as its motivating force.

This paper is structured as follows: Section 5.1. Section 5.2 provides an overview of the data and descriptive statistics, section 5.3 presents the identification strategy and section 5.4 provides the main results of the effect of the 3G upgrade on the number of reviews that restaurants receive on the Qype platform. Section 5.5 shifts the focus to the effect of mobile reviews on the rating of restaurants and section 5.6 shows a number of robustness tests. Section 5.7 concludes.

## 5.1 Related Literature

This paper contributes to the literature on the effects of infrastructure upgrades, especially on the effects of upgrades in telecommunications infrastructure. An early and influential empirical study on the returns to public infrastructure investments is Aschauer (1989). The study also initiated a debate about the problems of common time trends and reverse causality when estimating their effect on growth. Later studies show that when these issues are accounted for, the positive effects found by Aschauer almost vanish.<sup>5</sup> Therefore, policy makers should carefully evaluate which infrastructure investments are worth to be subsidized.

An important question in this context is, if infrastructure investments have positive returns in areas with low demand for the new infrastructure. While Hornung (2012) finds that Prussian cities that were connected to railroads for exogenous reasons experienced higher growth rates, Banerjee, Duflo and Qian (2012) find that new transportation infrastructure in China had no effect on per capita GDP growth for regions that were quasi-randomly connected to the transportation network. Duranton and Turner (2011) provide evidence that increases in lane kilometers of interstate highways positively affect the vehicle-kilometers traveled. Even for classic examples of public infrastructure investments, like the investments into the U.S. railroad network in the nineteenth century, the economic returns are highly controversial. While Jenks (1944) and a number of other early studies highlight the importance of railroads for the economic development of the USA, Fogel (1962) famously argues that the century-old river network in the United States was more effective for economic development than the heavily subsidized railroad network. Atack et al. (2010) finds that the railroad expansion into the Midwest had a significant effect on urbanization, but not on the growth of population density. Similarly, Fernald (1999) finds that large road infrastructure investments in the United States led to a one-time productivity boost in the 1950s and 1960s. On the margin, however, he does not find that road infrastructure investments are especially productive.

---

<sup>5</sup> For an overview of the discussion see Munnell (1992).



The literature on upgrades in telecommunications infrastructure is also anything but concordant when it comes to measuring economic effects. Waverman and Röller (2001) were some of the first to estimate the causal effect of telecommunication infrastructure investment on growth. Using data for 21 OECD countries for the period from 1970 to 1990, they find a significant positive effect that is especially strong when there is a high penetration of the technology. Czernich et al. (2011) also use data from OECD countries, but analyze them for the period from 1996 to 2007. Using an instrumental variable approach, they find that an increase in broadband Internet penetration has a positive effect on annual per capita growth. Kolko (2012) finds that the roll-out of broadband Internet increases population and employment growth, but does not affect wages or unemployment rates. Forman, Goldfarb and Greenstein (2012) show that investments into telecommunication technology are only associated with higher wages and employment for a small number of U.S. counties that were already well off prior to the investments. Czernich (2012) does not find any effects of broadband Internet infrastructure on unemployment rates in Germany.

Regardless of whether railroad, road or telecommunication infrastructures are concerned, the most important prerequisite for positive economic returns is that the upgraded infrastructure leads to a change in its usage. If the new infrastructure only changes usage in areas where the demand for the infrastructure was high anyway, one could argue that there is no use in subsidizing infrastructure investments in areas with low demand. This gives rise to the chicken-and-egg question of infrastructure investments: Will the availability of a new infrastructure stimulate demand or is demand a prerequisite for new infrastructure to have any effect? By using spatially fine-grained data, this paper finds that an upgrade in mobile Internet infrastructure leads to a change in its usage, even in rural areas where initial demand was likely to be low.

Given its focus on restaurant reviews, this paper is also connected to the growing empirical literature on online reputation and social learning. Many early empirical studies on the role of reputation use data from online auction platforms like eBay. Bajari and Hortacsu (2004) provide an overview of early studies that mostly use hedonic regressions to estimate the value of a good seller's reputation. Jin and Kato (2006) are one of the first to conduct an experiment on eBay and show that sellers who claim to provide extraordinary quality get paid more but do not actually deliver higher quality than the average seller. Resnick et al. (2006) estimate in an experiment that higher reputation increases willingness-to-pay. Later studies confirm this positive relationship between reputation and price (Lucking-Reiley et al., 2007) and highlight the especially detrimental effects of negative reviews (Cabral and Hortacsu, 2010).

The effect of reputation on sales has also been empirically explored for movies (e.g. Chintagunta, Gopinath and Venkataraman, 2010; Moretti, 2011), books (e.g. Chevalier and Mayzlin, 2006; Berger, Sorensen and Rasmussen, 2010) and many other experience goods (e.g. Hilger, Rafert and Villas-Boas, 2011; Luca and Smith, 2011). Recently, many studies have focussed on the effect of reputation reviews in the travel industry. Vermeulen and Seegers (2009) conduct an experiment in which they confront subjects with different online reviews of hotels and show that both positive and negative reviews raise the awareness for a hotel. Ye et al. (2011) explore the relationship between the number of reviews and the rating for hotels on the booking platform of a large Chinese online travel agency. By exploiting differences between two hotel review platforms, Mayzlin, Dover and Chevalier (2012) show that review manipulation occurs more for hotels with higher incentives to manipulate.

There are a number of studies that investigate review data for restaurants and are therefore closely related to this paper. Anderson and Magruder (2012) use a regression discontinuity (RD) design with Yelp data to show that a 0.5 star better rating leads to restaurants selling out 19 percentage points more frequently. Similar to Anderson and Magruder, Luca (2011) also uses an RD design and finds that a better star rating on Yelp increases revenues, especially for independent restaurants. Realizing the shortcomings of aggregated ratings which create discontinuities and other particularities, Dai et al. (2012) demonstrate how the average rating on restaurant review platforms should ideally be constructed and by how much the aggregation mechanism on Yelp deviates from this ideal design. Within this realm, Luca (2013) investigates the magnitude of review manipulation on Yelp and which restaurants are most likely to commit review fraud. Kang et al. (2013) use data from the Department of Public Health and from Yelp to show that restaurants with bad hygiene status are more likely to attract bad reviews.

## 5.2 Data

To estimate the effect of the upgrade from 2G to 3G on usage of the local online service Qype, I use data from five different sources. This section provides a detailed description of the data on restaurants and reviews as well as an overview of the spatial data used for identification. Section 5.2.3 provides summary statistics.

### 5.2.1 Restaurant and Review Data from Qype

I use a dataset of 105,954 restaurants in Germany that are listed on Qype. Qype was founded in 2006 and was Europe's largest user-generated reviews platform. According to the company, in 2012 the platform was used by 25 million unique users every month and more than 2 million users actively contributed by writing reviews.<sup>6</sup> The Qype app for the iPhone was consistently ranked as one of the top apps in the travel category of the Apple Appstore<sup>7</sup> in Germany. Restaurant reviews from the Qype platform were also used by third-party applications like Nokia Ovi Maps or Apple Maps. In October 2012, Qype was acquired by its larger US competitor Yelp.<sup>8</sup> All users and reviews from Qype were migrated to Yelp after the acquisition.

To get a sense of how location-based review platforms like Qype work, it is useful to take a look at how users experience them on mobile devices. Figure 5.1 shows screenshots from the Qype iPhone app. Apps for other operating systems, such as Android or Blackberry, work in a very similar way. In a first step, the user chooses what to search for. In addition to restaurant reviews, Qype offers customer reviews of bars, cafes, nightclubs and other businesses. In the first screenshot of Figure 5.1, the blue dot represents the user's current position and the red pins represent restaurants near to the user. When one of the red pins is tapped, the name of the respective restaurant is shown. When it is tapped again, the app presents the second screenshot in Figure 5.1. From here, the user can access detailed information such as photos, the restaurant's address, its average user rating, its phone number, and its web address. The third screenshot of Figure 5.1 shows all restaurant reviews that were written by previous customers. In addition to the text of the review, the app displays the star rating received by the restaurant from the review's author.

In addition to the information from the screenshots, I use many other restaurant characteristics, such as type of cuisine, possibility of booking a table online, and whether the restaurant actively promotes itself on Qype. Since every restaurant needs to be geocoded, I remove 1,778 restaurants with invalid addresses from the sample. Figure 5.2 shows the spatial distribution of the sample of restaurants. It is not surprising that there are fewer restaurants in the less densely populated eastern part of Germany, and there is, of course, a certain amount of clustering of restaurants in cities.

I combine this restaurant dataset with a database of 516,716 Qype customer reviews for all restaurants in Germany. Out of the 104,176 restaurants with valid addresses, 67,693 received at

<sup>6</sup> <http://de.press.qype.com/2012/09/10/eine-neue-dimension-fur-unternehmen-auf-qype/>

<sup>7</sup> <http://www.appannie.com/app/ios/299229792/ranking>

<sup>8</sup> <http://officialblog.yelp.com/2012/10/welcoming-qype-to-the-team.html>

least one review on Qype. On average, every restaurant in the sample received 4.96 reviews. Of those restaurants that received at least one review, the average is 7.63 reviews. As there is still a great deal of heterogeneity between the restaurants in the sample, I restrict the sample to those restaurants that had at least one review in 2008. This restriction is very important as the year 2008 marked the beginning of the smartphone app and mobile Internet hype. Apple's Appstore for the iPhone was introduced in July 2008<sup>9</sup> and the first phone with the Android operating system was introduced three months later.<sup>10</sup> Before 2008, third-party apps were available for very few mobile phones and their installation was quite complicated. By restricting the sample to restaurants that had at least one review in 2008, I ensure that all 25,340 restaurants in the sample had customers who were using Qype in the early days of the service. Restaurants that are only frequented by non-adopters are not considered in the remainder of my analysis.

## 5.2.2 Spatial and Regional Data

To obtain information about 3G coverage at a specific location, I use a unique dataset of all cellular antennas in Germany. This dataset includes 510,210 antennas on 63,481 antenna masts. Every antenna mast has several antennas, even if there is only one technology installed by one provider. For example, in order to cover a specific area with 3G, a network provider would need to install several antennas aimed in different directions on the same mast so that a large area is covered by the signal. In Germany, network providers often use their own masts and collocation of several providers on the same antenna mast is not as wide spread as in other countries. Figure 5.3 shows the spatial distribution of antenna locations in Germany, separated by the age of the antenna site. As expected, antenna density is highest in large cities, but because network providers need to cover even very remote areas, there is virtually no region without antennas. There is no obvious clustering of old compared to new antenna locations.

As the variation I am going to exploit for identifying the causal effect of 3G coverage is based on absorption of the 3G signal by obstacles like hills or forests, I also use a digital surface model (DSM). A surface model is a digital representation of the world's surface, including all objects on it. With the help of a DSM it is possible to determine the height of the surface at any point on the globe. One of the most advanced DSMs is the Shuttle Radar Topography Mission (SRTM), which dates back to a space mission in 2000. During the mission, the Space Shuttle Endeavour used a 60-meter-long mast and two special radar systems to scan the whole world from space (Farr et al.,

<sup>9</sup> <http://pogue.blogs.nytimes.com/2008/07/11/a-first-look-at-the-iphone-apps-store/>

<sup>10</sup> [http://www.gsmarena.com/t\\_mobile\\_g1-2533.php](http://www.gsmarena.com/t_mobile_g1-2533.php)

2007). For every arcsecond, which corresponds to an area of 30 square meters, the radar determined the height of the surface. As the raw data contain some voids and some empty cells, I use the SRTM data prepared by the CGIAR Consortium for Spatial Information (Jarvis et al., 2008). The data has a resolution of 3 arcseconds which corresponds to cells of approximately 90 square meters. Figure 5.4 shows a map section of the SRTM data for an area in southern Germany. The heights of the surface are visualized in a three-dimensional way. On the map, antenna locations are marked with red dots.

A special focus of this paper is on the effect of an upgrade to 3G in rural areas. Using administrative data at the county or municipality level to determine the rurality of a location would be problematic for this study. Larger counties with densely populated centers but vast areas of underpopulated land are bound to be misleadingly classified when administrative borders are used as a reference system. Therefore, I use a much more finely-grained population density measure. To determine whether a restaurant is located in a populated area, which might also be the center of a less densely populated municipality, I use data from a detailed population density grid created by the Joint Research Center (JRC) of the European Commission. The grid was created by desegregating administrative population data at the municipality level with the help of the CORINE Land Cover map, which provides information on whether a given square kilometer of the surface is populated or not (Gallego, 2010). By interpolation, the population density of a municipality is then broken down into a 1-hectare grid. Figure 5.5 depicts the population density in the area surrounding Munich. I define areas with a population density of less than 2,500 inhabitants per square kilometer as rural. These areas are marked as dark and grass green in Figure 5.5. Areas that are marked in shades of red, orange, and yellow are defined as urban. In Section 5.6, I will show that my results are robust to different definitions of rurality.

As figure 5.5 shows, many areas commonly classified as rural according to administrative data have relatively densely populated centers that fall into the urban classification used in this paper. This is the case because administrative boundaries are wider in rural areas and when the population is averaged across a large area of land, its density is considerably underestimated in the center. For mobile network operators, these urban centers of rural areas can be profitable since the installation cost per user can be as low as in cities if one antenna mast covers enough of the population. Conversely, the outskirts of cities are often classified as highly urban according to administrative data, but are more likely to be classified as rural according to the definition used in this paper. For mobile network operators, antenna locations are chosen based on the population that can be reached in the immediate surroundings of the antenna. If few people live on the outskirts

and the city center is already well covered by other antennas, operators may find it less profitable to install antennas in the outskirts area. The fine-grained population density data are therefore appropriate for the analyses in this paper.

I also use some additional data on the municipality level. From the Breitbandatlas Deutschland, I merge information on the diffusion of several Internet access technologies. The Breitbandatlas is an annual survey conducted by the German Ministry of Economics and Technology. All German mobile and fixed-line Internet providers are asked how much of the population in a specific municipality they cover with different broadband Internet technologies. The information from all Internet providers is then aggregated and the resulting dataset informs about the share of inhabitants in a municipality that has access to a certain Internet access technology. For this study, I am mainly interested in DSL coverage for the year 2008. I also merge information from the German Federal Statistical Office on the population and the area of every municipality.

### **5.2.3 Descriptive Statistics**

Table 5.1 presents some descriptive statistics for the data on the restaurant level. Mean, standard deviation and the number of observations are shown for restaurants in the full and rural samples. The first set of variables provides some basic information about review characteristics. On average, restaurants have more than 13 reviews in the full sample and about 8 reviews in the rural sample. This difference between samples is also apparent in 2008, although in absolute levels, it only amounts to 0.8 reviews. The statistics for mobile reviews and reviews from different devices all refer to the post-treatment period 2012. Row 4 and 5 show that the overwhelming majority of reviews stems from the iPhone. In both samples, the share of iPhone reviews on all mobile reviews is 76 percent. The share of mobile reviews on all reviews is also equal in both sample and amounts to 13 percent. As described above, I restrict both sample to restaurants that have at least one review in 2008. This leads to the average restaurant in the full sample having its first review in April 2007 and the average restaurant in the rural sample having its first review at the end of May 2007. The author experience variable is a measure of how many other reviews are written by the authors of reviews for a particular restaurant. This variable should capture how many reviews are submitted by authors who only reviewed one restaurant and how many reviews are submitted by experienced authors. The former group of reviewers is more likely to be biased or even connected to the restaurant owner, while the latter group potentially has a reputation to lose.

The second set of variables in Table 5.1 shows characteristics of the mobile network around restaurants. In the full sample, 75 percent of restaurants have a clear line of sight to an antenna, in the rural sample the share amounts to only 59 percent. Details on why and how the clear line of sight between restaurants and antennas is calculated, are provided in Section 5.4. The average number of antennas is also higher in the full sample. The variables in third and fourth row are proxies for the density of the mobile network around restaurants. In the full sample, there are on average 422 antennas, in the rural area there are on average 177 antennas within a radius of 10 km. Only these antennas are considered for the calculation of the lines of sight between restaurant and antenna, because antennas that are further away than 10 km are very unlikely to provide a strong enough 3G signal. Both in urban and in rural areas, the number of antenna locations is inflated by the number of network providers. In Germany, network providers mostly have their own antenna sites where only antennas of one provider are installed. Given that there are three major mobile network providers, it is not uncommon that several antennas sites are considerably closer to each other than it would be necessary if there was only one provider. The differences between the full and the rural sample of restaurants can be explained by the number of users that share one antenna. Every antenna cell has a maximum number of users which share the capacity of the cell. When the maximum number of users is exceeded or a higher capacity per user is demanded, the network provider has to install further antennas even though the existing antennas could theoretically cover the area of a new antenna site.

The third set of variables in Table 5.1 present regional characteristics of the municipality and immediate area around restaurants. It is not a surprise that population, area and population density are all lower in the rural sample, compared to the full sample. After all, the rural sample is restricted in a way that less populated municipalities are overrepresented. My definition of rural areas is based on the density cell variable in row four of the regional variables. In the full sample, the average population density cell is above 5000 inhabitants. Among those restaurants in cells below 2500 inhabitants, which is about half of the average in the full sample and the condition which I use to restrict the rural sample, restaurants are in a 636 inhabitants cell, on average. I am also able to control for fixed-line broadband coverage on the municipality level, denoted by the DSL variable. While in the full sample, fixed-line broadband coverage is above 90 percent, it is 5 percentage points lower in the rural sample.

The last set of variables describes restaurant characteristics obtained from the Qype website. The first variable shows that 24 percent of restaurant owners in the full sample have "claimed" the profile of their restaurant on Qype, meaning that they identified themselves and provided additional

information. 67 percent of restaurants in the full sample and 62 percent of restaurants in the rural sample have at least one photo on their profile. Photos can be either uploaded by restaurant owners themselves, or by customers who leave a review. The next two variables show the percentage of restaurants that have a phone number and web address on their profile. 6 percent of restaurants in the full and 5 percent of restaurants in the rural sample offer the possibility to book a table online. The rank variable indicates how high the restaurant ranks within its municipality. Since there are on average fewer restaurants per municipality in the rural sample, the average rank is also halve of that in the full sample. Qype users are also able to use the smartphone app in order to "checkin" into a restaurant. By checking in everytime they visit a restaurant, they show other users that they are regular visitors of the place. The last four variables are examples of the more than 80 restaurant dummies that describe the cuisine or type of a restaurant. While in the full sample, 20 percent of restaurants have German and 17 percent Italian cuisine, in the full sample the share of German restaurants is 31 and the share of Italian restaurants only 10 percent. Differences in the shares of Chinese and Greek restaurants are less pronounced.

### **5.3 Identification Strategy**

This paper investigates the relationship between an upgrade in the mobile Internet infrastructure and usage of the local online service Qype. To identify the causal effect of 3G coverage, I estimate a difference-in-difference model where I compare restaurants with and without 3G reception before and after mobile Internet became a mass phenomenon. The difference-in-difference setting has the advantage that I can avoid bias from permanent differences between restaurants with and without 3G and also exclude any bias from time trends, such as the increasing popularity of online services over time. I can further control for a wide set of regional and restaurant characteristics in order to ensure that there are no observable differences in characteristics that could drive the trends of the outcome variables.

The key remaining assumption of the difference-in-difference model is that restaurants with and without 3G have common underlying trends in the number of reviews they receive on Qype. I try to reduce the risk of differences in underlying trends by exploiting exogenous geographic variation for separating the treatment from the control group. Instead of using actual 3G coverage of a restaurant, a measure that is not even available at such a fine-grained level, I use a variable that indicates if there is a clear line-of-sight between the immediate surrounding of a restaurant and mobile network antennas in that area. Restaurants that do have a clear line-of-sight to at least one



mobile network antenna within 10 kilometers are considered treated, while restaurants that do not have a clear-line-of-sight are considered untreated.

The reasoning behind this rather complicated strategy is that given the antenna density in an area and the distance to the next antenna, a clear line-of-sight was not very important when 2G was the prevailing wireless technology. In other words, making phone calls over the 2G network was possible even if there were obstacles between the phone and the next antenna. For 3G, which uses a much higher frequency band in Germany, obstacles between phone and antenna suddenly became much more important. One way to experience this first hand, is to go inside a large building or into the basement of a building and observe the network reception of the mobile phone. The 3G signal typically becomes weaker and at some point the phone will switch to 2G. But just because 3G is not available inside a building does not mean that phone calls are not possible. Often the 2G network has still quite impressive reception when 3G is completely blocked by the walls of the building. This difference in the absorption of 3G compared to 2G signals is exactly what I am exploiting in this study.

Using the exogenous variation to separate treatment and control groups can be compared to a reduced-form estimation in an instrumental variable setting. Since the actual 3G signal strength at a restaurant cannot be observed, a conventional IV approach is not viable. A simple regression of the actual 3G availability on the number of reviews would most likely overestimate the effect of the 3G upgrade. Areas with higher growth trends can be assumed to also have a higher smartphone penetration and could therefore get preferred treatment by the network providers. These unobserved characteristics are also likely to be correlated with the number of reviews that restaurants receive. The "reduced-form" estimation comes at the cost of introducing noise into the measurement of the effect. However, it is unclear whether this noise from the geographic variation will bias the results upwards or downwards. Having a clear line of sight does not strictly correlate with most other geographic measures, such as the slope of the terrain. Indeed, both very high terrain ruggedness and very low ruggedness reduce the chances of a clear line of sight. Nolen and Klonner (2010) show that the best conditions for a mobile network are relatively flat areas with hills in between. When antennas can be placed on top of hills, it is most likely that the flatter areas surrounding the hill will have a clear line of sight. As the geographic variation I use is very specific, I argue that there is no obvious under- or overestimation of its effect on reviews.

Figure 5.6 shows schematically what the treatment variable measures. Restaurants A and B do both have the same distance to an antenna mast and the number of antenna masts around the two

restaurants is the same. Both restaurants are well covered by 2G. While restaurant A has a clear line-of-sight to the next antenna, for restaurant B there are forests or hills between the restaurant and the antenna. The latter is not a huge problem for the 2G signal, as it will pass around the hills and through or above the forests, similar to a radio signal that passes through and around many different kind of obstacles until it reaches a radio receiver. For the 3G signal on the higher frequency band, depicted by the line with shorter breaks in Figure 5.6, the obstacles between the antenna and restaurant B constitute a larger problem. Unlike the 2G signal, it will most likely be blocked by any major obstacle on the way and will therefore not reach restaurant B. Restaurant A, albeit being equally far away from the antenna, will be covered by 3G.

The way I apply this identification strategy to the data is to draw lines between every restaurant and all surrounding antenna locations. Then I use the digital surface model to determine the height of the surface in the immediate area surrounding the restaurant and the immediate area surrounding the antenna mast. In a computationally very demanding process, I then determine if there are any major obstacles, measured as elevations on the surface, along the line between the restaurant and the antenna. For simplicity, in the remainder of this paper I will call an antenna visible if there is no major obstacle between the antenna and the restaurant. After determining the visibility of every antenna around the restaurant, I aggregate the information and create the treatment dummy which is one if there is at least one visible antenna and zero if none of the surrounding antennas are visible.

The treatment dummy creates the first difference I am using, namely the difference between restaurants that are likely to have and restaurants that are likely to not have 3G due to geographic variation. The second difference I am using is the difference between the time when 3G and with it mobile broadband Internet was not used to a meaningful extent and the time when 3G was a mass market phenomenon. There are a couple reasons why I decide to not choose the year of the 3G spectrum auction, 2000, or the year when the first 3G network went online, 2004, as a base year. First, few online services that are based on location data and that were later used as smartphone applications existed back in 2000 or 2004. The location aware Internet is a fairly recent phenomenon. It would therefore be hard to compare a period without any comparable service to a period where smartphones and mobile Internet are ubiquitous. Second, as it is often the case with new infrastructures, there are important complementarities that have to be taken into account. In the case of 3G, the obvious complementarity is the smartphone. Before 3G enabled smartphones with user interfaces that are build for browsing the Internet were on the market, even the best 3G network was of little value to consumers. I therefore want to choose a point in time when location

based services like Qype were already existent but the 3G network, and with it mobile Internet as we know it today, was not yet popular due to the lack of complimentary products.

For many reasons, 2008 was the year in which 3G started to become a technology for the masses. Two corporate events fueled the popularity of 3G. Most importantly, in July 2008, Apple introduced its iPhone 3G, arguably the device that contributed the most to the mobile Internet revolution. In the same month, Apple opened its App Store, which allowed users to install third-party applications. Before Apple revolutionized the way online content is consumed on mobile phones, many in the industry believed a "new version" of the Internet was needed before a breakthrough in mobile Internet would be possible (West and Mace, 2010). After Apple's success, the competition followed suit and the first Android phone was introduced in October 2008. In the same month, Google opened a competing app store, the Android Market. The combination of modern devices like the iPhone, smartphone apps that deliver online content optimized for small touch screens, and the 3G broadband network led to a growing popularity of mobile Internet. At the same time, location based online services like Qype realized the opportunities of mobile broadband Internet and started to develop apps that allowed the use of their services on the go. As 2008 marks the year of the first mass market 3G smartphones and also the year when existent location based services started to focus on mobile Internet users, I choose 2008 as the base year of the difference-in-difference estimations.

Figure 5.7 shows the data volumes in the German mobile network and confirms that before 2008 there was very little data traffic. This reflects that before 2008, complimentary products for utilizing the 3G network were rare. Although it is possible to use mobile Internet on 2G devices, the experience is very different, as the speed of the Internet is only a fraction of the Internet speed in 3G networks. Loading pictures or maps in the 2G network is often a very time-consuming undertaking. Nevertheless, Figure 5.7 depicts a small increase of data volumes in 2008. I therefore show in Section 5.6 that my results are robust to choosing 2007 as the pre-treatment period. For all other estimations, 2008 is considered pre-treatment. Since I am looking at the cumulative number of reviews as an outcome variable, I choose 2012 as the second point in time that I compare the pre-treatment period with. Later points in time could be affected by the next generation of mobile network technology, 4G.

Figure 5.8 shows the number of total reviews and the number of reviews from mobile devices for treated restaurants (with visible antennas) and for untreated restaurants (without visible antennas). While the two graphs on the left show the development of reviews for the full sample of restaurants,

the two graphs on the right show the development for restaurants in rural areas. For the total number of reviews, there is a parallel trend for restaurants with and without antenna visibility until 2008. After 2008, the trends for the two groups clearly diverge and restaurants with visibility experience a stronger increase in reviews than restaurants without visibility. Although the levels are considerably smaller in the rural sample, the trends are similar to the full sample. I do not observe mobile reviews in 2008. When the first Qype app was released in December 2008, it allowed to search for nearby restaurants, but did not allow for the submission of reviews from the mobile phone. Due to the early focus on search, the app was called "Qype Radar". In April 2009, the company released an update for the app, both for iPhones and Android devices, that allowed to post reviews directly from the app<sup>11</sup>. The two lower graphs in Figure 5.8 show that until 2009, the trends of mobile reviews are parallel for restaurants with and without visibility, even though the absolute levels are very low. After 2009 the trends of mobile reviews diverge between the two groups.

I estimate the following difference-in-difference model in which I compare restaurants with and without visibility in 2008 and 2012:

$$y_{it} = \beta_0 + \beta_1 \text{visibility}_i + \beta_2 2012_t + \beta_3 (\text{visibility}_i \times 2012_t) + \beta_4 X_i + \beta_5 R_i + \beta_6 \text{type}_i + \beta_7 \text{county}_i + \varepsilon_i$$

where  $y_{it}$  is the number of reviews that a restaurant  $i$  receives in period  $t$ .  $\text{visibility}_i$  is a dummy that takes the value 1 if there is a clear line-of-sight between the area surrounding restaurant  $i$  and any antenna mast within a radius of 10 km. Antennas at a distance of more than 10 km are highly unlikely to provide a strong enough 3G signal.  $2012_t$  is a dummy that is 1 in the post-treatment period 2012 and  $\text{visibility}_i \times 2012_t$  is the interaction term for visibility in 2012, meaning that it becomes 1 for having a clear line-of-sight in the post-treatment period.

Vector  $X_i$  includes a set of characteristics of restaurant  $i$ , namely dummies that indicate whether the restaurant actively uses Qype for promotional purposes, whether there are photos of the restaurant available on Qype, whether it is possible to book a table at the restaurant online, whether there is a link on Qype to the restaurant's website, and whether there is a phone number for the restaurant available on Qype.  $X_i$  also includes the average star rating of restaurant  $i$ , its Qype ranking within its municipality, the year in which the restaurant received its first review, and the accumulated experience of all authors who wrote reviews for the restaurant. The latter is measured by the sum of all reviews that were written by the respective author.

<sup>11</sup> <http://www.theguardian.com/media/pda/2009/apr/09/mobilephones-iphone>

Vector  $R_i$  includes regional characteristics such as the population and area of the municipality in which restaurant  $i$  is located. Very importantly,  $R_i$  contains controls for the distance to the next antenna mast as well as the number of antennas within 10 km around the restaurant. While the latter is a good proxy for the mobile network density in the area, the former controls for the remoteness of the restaurant relative to the mobile network. Given the distance and the density of antennas, a clear line-of-sight is the main determinant of 3G reception.  $R_i$  also contains the average fixed-line broadband Internet penetration in the municipality of restaurant  $i$ .  $type_i$  is a set of more than 80 dummies that indicate the kind of cuisine or if cuisine is not applicable, the kind of establishment, for example beer garden or fast food restaurant.  $county_i$  is a set of county-fixed effects.

## 5.4 Effects on the Number of Restaurant Reviews

Table 5.2 presents results for the difference-in-difference estimation described above with the total number of reviews, both mobile and non-mobile, as the dependent variable. The first four columns show estimates for the full sample of restaurants that had at least one review in 2008. The interaction term in the first line shows that those restaurants with a clear line of sight received significantly more reviews than restaurants without a clear line of sight, even after controlling for differences in levels and a common underlying trend. Having a clear line-of-sight increases the number of reviews by about 3. The time trend itself also has a strong effect on the number of reviews which is not surprising, given the increasing popularity of local online services like Qype and the fact that the number of reviews are cumulatively measured. Columns 5 to 8 show estimations for the reduced sample of restaurants in areas that I define as rural. The point estimates in the rural sample are significantly smaller than in the full sample which partly reflects that the average number of reviews is considerably smaller in rural areas, as shown in Table 5.1.

It is striking how little the coefficients change when restaurant characteristics, restaurant type-fixed effects, regional characteristics and county-fixed effects are included in the estimations. Also restaurant-fixed effects do not change the coefficients in columns 4 and 8 by a lot. This indicates that the distributions of restaurant and regional characteristics in the two groups are either very similar or that they do not have any distinct effect on the trends of restaurant reviews. If they had any effect on the level of reviews, this could not be observed in the estimations, as the difference-in-difference model already accounts for any differences in levels between the two groups.

In Table 5.2, the dependent variable is the total number of reviews which is not only driven by reviews submitted by mobile phones but also - and in fact to a larger extent - by reviews submitted by PCs. There are several reasons why 3G availability also has an effect on non-mobile reviews. First, the early focus of the Qype app was discovery of nearby restaurants. This is also reflected by the name of the app which was "Qype Radar" until the end of 2010. The aim was to provide the user with a "radar" for nearby restaurants and show the average rating as well as reviews of the restaurant. At this time, many users probably used the app to discover a restaurant but used their PC at home to write a review after the visit. Second, even after the app had a stronger focus on the rating of restaurants, many users may have disliked to write a review on the phone, for example because of the smaller keyboard. Third, the information on restaurants in Qype was not only available through the Qype app. As early as May 2010, Qype started a cooperation with Nokia so that restaurant reviews were henceforward displayed in the Nokia Ovi Maps app.<sup>12</sup> A similar partnership followed later with Apple Maps, the default maps application on iPhones. However, these third-party apps only displayed the rating and reviews of restaurants but were not able to directly submit a new review. It is therefore reasonable to assume that 3G also affects non-mobile reviews by increasing the mobile discovery of restaurants through information from the Qype platform.

Although there are good reasons to believe that non-mobile reviews are affected by 3G, the more obvious question is if 3G availability increases the probability that a restaurant receives reviews which are submitted from mobile phones. This corresponds to the scenario that a user discovers a restaurant through the app and opens the app immediately after his visit to the restaurant in order to leave a review. Table 5.3 shows that the probability for a restaurant to receive at least one mobile review increases by about 9 percent when the restaurant has a clear line-of-sight to an antenna. For restaurants in rural areas, the effect is almost the same as in the full sample of restaurants. Adding controls on the regional and restaurant level as well as adding restaurant-fixed effects does not change the coefficients. The visibility dummy in the third row has coefficients that are close to zero as both groups, restaurants with and without visibility, do not have mobile reviews in the pre-treatment period. The difference-in-difference model therefore only accounts for the trend in the control group, but does not have to account for differences in the levels between groups, as they are not existent.

One potential problem of the estimates in Tables 5.2 and 5.3 is that restaurants in the treatment and control groups may have specific Internet related trends. Let us assume, for example, that a restaurant owner discovers the revenue-enhancing potential of the Internet, sets up a website, and

---

<sup>12</sup> <http://netzwertig.com/2010/05/28/linkwertig-qype-google-buzz-bbc-youtube/>

starts advertising on Facebook and Google. These activities might induce some Internet users to visit the restaurant and subsequently leave a review on Qype. If owners of restaurants with 3G coverage are more likely to discover the potential of the Internet, this could bias the estimates of the effect on the total number of reviews. One way to reduce this risk of potential bias is to estimate the effect on the share of mobile reviews over all reviews. If restaurant owners in areas with 3G are for some reason especially enthusiastic about the Internet, this could have an effect on the trend of mobile and on non-mobile reviews. However, it is much harder to think of a reason why it should affect the trend of mobile reviews relative to non-mobile reviews. By estimating the effect on mobile over non-mobile reviews, I implicitly control for something like restaurant-specific Internet trends. In order to bias these estimates, there would have to be differences in the trends of mobile reviews that do not affect the trends of non-mobile reviews. Besides mobile Internet availability itself, it seems unlikely that there are many factors driving these relative trends.

Table 5.4 presents estimates with the share of mobile over non-mobile reviews as the dependent variable. The first four columns show that there is almost no effect on the share of mobile over non-mobile reviews in the full sample. This is surprising and could either indicate that the increased discovery of restaurants through the Qype app affects mobile and non-mobile reviews to a similar extent, or it could indicate that there are restaurant-specific Internet trends other than 3G which drive the results in the previous tables. Columns 5 to 8 show that in the rural sample, restaurants with a clear line-of-sight have a robust 2 percentage points higher share of mobile reviews. Given that in the rural sample the aggregate share of mobile reviews is below 16 percent, an increase by 2 percentage points is clearly noticeable. One possible explanation for the larger effect in the rural sample is that Qype users in rural areas rely more on the mobile phone. Maybe, in urban areas Qype is more often used to discover new restaurants, while in rural areas it is more often used to discover a restaurant at all. If users who find restaurants with Qype in rural areas are more often travelers that do not have a PC available, they might prefer to submit a review on the mobile phone instead of not submitting a review at all. A different explanation could be that the aforementioned third-party apps that use Qype data but do not allow for the submission of reviews are more frequently used in urban than in rural areas.

To summarize the main results, Tables 5.2, 5.3 and 5.4 show that restaurants with a clear line-of-sight get about 3 more mobile and non-mobile reviews, a number which is smaller in rural areas due to the lower absolute number of reviews. The probability for restaurants with visibility to get at least one mobile review is around 9 percent higher and the effects are similar for restaurants in urban and rural areas. Concerning the relative increase of mobile reviews compared to non-mobile

reviews, I only find a significant increase by 2 percentage points for restaurants in rural areas. I interpret the effect of having a clear line-of-sight on the number of reviews as the effect of the increased likelihood to have 3G reception in the area around the restaurant. The higher number of reviews indicates that more customers use Qype to discover nearby restaurants when 3G is available. Both the increased discovery and the increased number of reviews are signs that the upgrade to 3G has led to the adoption of location-based online services and has induced a change in consumer behavior.

## 5.5 Mobile Reviews and the Restaurant Rating

In section 5.4, I have seen that 3G availability increases the total number of reviews and the share of mobile reviews for a restaurant. This shows that the upgrade to 3G infrastructure had a direct effect on the uptake of mobile online services. Besides the more general question about the relationship between infrastructure upgrade and uptake, the specific application in this paper leads to the question if the uptake of mobile online services has any positive economic effects. From a welfare perspective, more restaurant reviews lead to more market transparency for consumers. Since a restaurant is an experience good, where consumers do not have perfect knowledge about product characteristics and prices, more transparency is usually thought to promote competition and increase consumer welfare (e.g. Schultz, 2004; Varian, 1980). According to Anderson and Magruder (2012), a larger number of reviews also reduces the likelihood that restaurant owners game the review system which is further increasing transparency.

While the positive effects of reviews on transparency and consumer welfare are apparent, it is much less obvious if the restaurants benefit from a larger number of reviews. On the one hand, the platform opens a new marketing channel for restaurants and more reviews are a good way for restaurants in remote areas to attract the attention of consumers. On the other hand, the increased transparency might harm restaurants that offer low quality or above market prices. It is also unclear whether restaurants benefit from the possibility that customers submit their review immediately from their mobile phone instead of several hours or days later from their PC at home. On the one hand, customers that are very happy about their experience might want to give an immediate positive review with their mobile phone and possibly forget about it when they are back at home. On the other hand, customers that had a disappointing experience at a restaurant might be more emotional or filled with anger when they write their review immediately. After a little while, their



anger might cool off and they could either decide to not leave a negative review at all, or write a less destructive review at their PC at home.

### 5.5.1 Do Mobile Reviewers Give Better Ratings?

In order to analyze if reviews that are submitted from a mobile phone are generally more or less generous than reviews that are submitted from a PC, we will shift the focus from the level of restaurants to the level of reviews. Table 5.5 shows the relationship between the star rating of a single review, ranging from 1 to 5, and the source of the review as well as other review characteristics. Note that the sample includes reviews from all restaurants and not only from restaurants that had at least one review in 2008. The first column shows that reviews submitted from a mobile phone have 0.11 more stars than reviews written on the PC. Although the coefficient is highly significant, its size is not overwhelming. A restaurant that only has PC reviews would need ten times as many mobile reviews to raise its average rating score by 0.1. In the second columns, additional controls on the review level are added to the regression. On Qype, users can comment on every review and the second row of Table 5.5 shows, that the number of comments is associated with a lower star rating. The reason for this could be that restaurant owners often use comments to react to reviews. When a review is bad, restaurant owners have a stronger incentive to react and explain why the reviewer has had an experience that is not representative for the restaurant. Users also have the opportunity to compliment reviews. Reviews that are especially long and well written usually get a large number of compliments. As shown in the third row of Table 5.5, the number of compliments is very weakly correlated with the rating. Finally, there is no evidence that review authors with larger experience, measured as the total number of reviews submitted, generally give better or worse ratings.

Columns 2-4 of Table 5.5 show that the correlation between the star rating and the source of a review hardly changes when restaurant controls, restaurant type dummies, regional characteristics or county-fixed effects are accounted for. In order to further exclude potential bias from omitted restaurant characteristics, I add restaurant-fixed effects in columns 5-6. In these estimations I do not find significant differences compared to the estimations without restaurant-fixed effects. This indicates that unobserved restaurant characteristics, like the quality of the food and competence of the waiters, do not confound with the correlation between the source and the rating of a review. If better restaurants were both more likely to receive a good rating and reviews submitted from mobile phones, I would expect the coefficients in the estimation with restaurant-fixed effects to be significantly smaller than the coefficients in columns 1-4.

Even after including restaurant-fixed effects, the small but positive relationship between mobile reviews and their rating could be driven by unobserved characteristics of the reviewer. If young reviewers gave better reviews and were also more likely to use their mobile phone to submit a review, my estimates could be upwards biased. By comparing mobile and PC reviews of the same author, I am able to hold all reviewer characteristics, like age, income and education, constant. I also account for the fact that some customers primarily leave reviews when they are satisfied and other customers primarily leave reviews in order to complain about a bad experience. In my sample, I can identify more than 150,000 authors that leave about 3 reviews on average. The distribution of reviews is positively skewed, so that many authors leave much more than the average number of reviews. However, the majority of authors either always uses the PC to leave a review or always uses the mobile phone. There are 8,660 authors who use both their PC and their mobile phone to leave a review. These authors submit 10.8 reviews on average, which is more than three times as many reviews than those submitted by the average author. Since I will only identify from this potentially special subsample of authors when I include author-fixed effects, I will restrict the sample to these authors in Table 5.6.

The first two columns of Table 5.6 show benchmark estimations with restaurant-fixed effects in the sample of authors who left at least one mobile and one PC review. Compared to the restaurant-fixed effects estimations in the full sample, the positive correlation between mobile reviews and their rating is much weaker in the restricted sample. When I include author-fixed effects in columns 3-6, the small positive coefficients turn negative but stay quite small in size. In the within-author estimations, the correlation between the rating and the number of comments and compliments becomes also very small. Including controls for restaurant characteristics and restaurant type dummies hardly changes the coefficients. This indicates that when unobserved author characteristics are held constant, mobile reviews are, if anything, associated with slightly less generous restaurant ratings. Since authors probably visit and review the same kind of restaurants when they go out, observable restaurant characteristics do not influence the results.

So far, I have looked on mobile reviews compared to PC reviews and accounted for specific characteristics of mobile reviewers by including author-fixed effects. However, I ignored that not all mobile reviewers are the same. Most importantly, there is a lot of evidence about the socioeconomic differences between users of different smartphone brands. For example, iPhone users have on average higher socioeconomic background than Android users, spend more money on Apps and

spend more time browsing the web with their phone.<sup>13</sup> Blackberry smartphones are more often used as business phones, especially by employees of large corporations and banks. The differences between smartphone users could lead to different correlations with the star rating of reviews. In Table 5.7, I therefore compare the relationship between reviews from different smartphone brands and their corresponding rating.

Columns 1-4 of Table 5.7 present results of within-restaurant estimations with smartphone brands entering the estimations consecutively and simultaneously. The results indicate that reviews from Blackberries are associated with a considerably higher rating than reviews from Android smartphones and iPhones. Among all three phone brands, iPhone reviews have the weakest positive correlation with rating. This pattern is confirmed by the within-author estimations in columns 1-4. In order to show from how many observations the identification stems, I restrict the sample to authors that submitted at least one review from the PC and one review from a smartphone of the respective brand. While this leaves me with quite a respectable sample size for iPhones and Android phones, the sample for Blackberries is reduced considerably to 3,910 reviews by 346 authors. Since most users only have one smartphone, the brand effects can not be estimated simultaneously for all three smartphones in the within-author specifications. For iPhones, I find a small but significant negative correlation with rating, similar in size to the within-author estimates of combined mobile reviews. For Android, the negative estimate is only half as large as for iPhones and for Blackberries the correlation is virtually zero.

Tables 5.5, 5.6 and 5.7 show that there is a very weak relationship between the source of a review and its rating. The positive sign of estimates in the within-restaurant estimations might indicate that given the quality of the restaurant, customers that use mobile phones are slightly more generous in their rating. The negative sign of estimates in the within-author estimation could indicate that given customer characteristics, reviews that are submitted from mobile phones tend to be a little bit less generous. Overall, the magnitude of the different correlations is, however, very small. It is unlikely that restaurant ratings are biased in either direction due to mobile reviews. In order to be sure about the effect, I need to turn back to the restaurant level and analyze if the number of mobile reviews has an effect on the average star rating of a restaurant.

---

<sup>13</sup> [http://www.slate.com/blogs/business\\_insider/2014/04/04/apple\\_vs\\_android\\_developers\\_see\\_a\\_socioeconomic\\_divide.html](http://www.slate.com/blogs/business_insider/2014/04/04/apple_vs_android_developers_see_a_socioeconomic_divide.html)

### 5.5.2 Are Restaurants with More Reviews Better Rated?

Table 5.8 presents effects of the number of PC and mobile reviews on the average rating of a restaurant. Columns 1-4 show that neither the number of mobile reviews, nor the number of PC reviews has effects on the rating. Due to the large sample size, the effects are estimated very precisely, so that most of the effects are statistically significant, although they are virtually zero. Including a set of controls on the restaurant level as well as county-fixed effects has very little effect on the coefficients in the pooled estimations of columns 1-4. In column 4, I add restaurant-fixed effects. I use variation in the number of reviews and the average rating from every year between 2008 and 2012. Holding unobserved restaurant characteristics constant, I still do not find any effect of the number of reviews. Columns 5-8 repeat the same estimations in the rural sample. Although the standard errors are considerably larger than in the full sample which can be explained by the smaller sample size, the coefficients of the number of mobile reviews do not differ between the two samples. For the number of PC reviews, the coefficients are larger in the rural sample, but still very small.

The results in Table 5.8 can be interpreted as evidence against a bias in restaurant ratings when the rating is based on a small number of reviews. Even though restaurants with fewer reviews face stronger incentives to game the system and fake reviews in order to improve their rating (Anderson and Magruder, 2012), this is apparently not what is happening to a large degree. If restaurants with a small number of reviews were more likely to game the system, we would expect a significant negative relationship between the average score and the number of reviews. Based on the coefficients from the estimation with restaurant-fixed effects, the average restaurant would need 278 more PC reviews or 182 more mobile reviews to reduce its average rating by one star. Since the average number of reviews for restaurants that have at least one review is about 7, it is unlikely that restaurant ratings are affected by the number of reviews in a meaningful manner.

Restaurants do not seem to benefit from more reviews in terms of a better rating. Nevertheless, the overwhelming majority of reviews contain information that goes beyond the simple 1-5 star rating. In my sample, only 10 percent of all reviews have less than 45 characters of text and more than 60 percent of reviews have at least 250 characters, corresponding to about 50 words. The text of a review often contains information about the price and the quality of the food or the friendliness and competence of the service. It may also contain recommendations regarding specific dishes or which kind of table one should ask for in the booking process. Many reviewers also upload photos of the restaurant or the food they had. Since these photos are typically taken with mobile

phones and uploaded through the app, they are more frequently available for restaurants that receive mobile reviews. This sort of information can have significant effects, as several previous studies have shown. For example, Cai, Chen and Fang (2009) find that providing restaurant visitors with information on the most popular dishes raises the demand for these dishes significantly. Chevalier and Mayzlin (2006) show that consumers respond to the text length of reviews rather than just on their rating. It is also quite obvious that this kind of information improves market transparency and should therefore increase competition and consumer welfare. But also from the perspective of good restaurants, the additional information conveyed by every review should be seen as beneficial, even though it does not improve the rating directly. For restaurants with good ratings, a larger number of reviews has the additional advantage that the last review has a higher probability of being fairly recent and potential customers are likely to trust recent ratings more than older reviews.

## 5.6 Robustness

Our identification method presented in Section 5.4 tackles a large number of possible sources for bias in my results. First of all, the difference-in-difference estimates account for different levels in reviews for restaurants with and without 3G. They also account for time trends that are induced by the increasing popularity of online services and other trends that affect both groups equally. In order to eliminate the risk that there are differences in the distribution of characteristics which could affect the trends between the treatment and control group, I include a large set of controls as well as restaurant-fixed effects in my estimations. Furthermore, I mitigate the risk of differences in underlying trends between treatment and control group by exploiting exogenous variation in 3G availability as treatment. This brings us into the direction of random assignment between treatment and control groups. My identification would be weakened if the variation I am using as treatment was not only based on exogenous geographic variation but also on endogenous decisions by network providers. In the following section I am trying to mitigate this risk. I also confirm that the results are robust to different sample restrictions.

### 5.6.1 Threats to the Exogeneity of a Clear Line-of-Sight

One reason why I use clear line-of-sight as a proxy for 3G reception is that given the distance to an antenna, a clear line-of-sight has a strong impact on 3G reception, but affects 2G reception to a much smaller degree. This leads to some exogenous variation in the availability of 3G. The 2G

antenna masts covered the whole of Germany reasonably well before the introduction of 3G. When 3G was introduced, network providers started to upgrade their 2G masts with the new technology. But even if all 2G masts were upgraded to 3G, there would still be many areas left without 3G coverage. This has two reasons. The first reason is that the higher frequency of the 3G signal leads to a smaller radius covered by each antenna. The second reason is that obstacles, like forests, hills or large buildings, absorb the 3G signal stronger than the 2G signal. In this paper, I only exploit the second property while I control for the first. However, the exogeneity of the second property, namely the different absorption properties of 2G and 3G, could be threatened by the strategic choice of locations for additional 3G antenna sites.

In 2001, few months after the first UMTS spectrum auction in Germany, there were around 40,000 antenna locations in Germany (Umweltausschuss des Bundestages, 2001). At this time, mobile network operators estimated that they would need around 15,000 new antenna sites for the first stage of the 3G roll-out in Germany. Mainly due to the smaller radius of the 3G signal, network operators planned to build new antenna masts from the beginning of the roll-out. It is likely that economic factors played a role in the selection of new antenna locations. Areas with larger population density were probably given priority, as in those areas more customers can be reached with less antennas. It is also likely that areas with higher purchasing power were given priority, because in these areas the likelihood of attracting customers for the more expensive 3G data plans was higher.

Let us assume that there are two restaurants which both have 2G reception, but not 3G reception due to the lack of a clear line-of-sight to the next antenna mast. One of the two restaurants is located in an area with strong purchasing power while the other is located in an area with low purchasing power. It is likely that network providers will first build antenna masts that cover the restaurant in the high purchasing power area. If this was the case, the exogenous variation in 3G reception would be disturbed by endogenous investment decisions of the network providers. This could lead to an overestimation of the effect of 3G reception on the number of reviews that a restaurant receives. Although the difference-in-difference estimations account for differences in pre-treatment levels, they assume that areas with and without visibility have a common pre-treatment trend. This pre-treatment trend is especially hard to confirm for mobile reviews as they could not have been submitted in the pre-treatment period. If network providers choose new antenna locations based on trends, for example in purchasing power, this would be a serious threat to identification.

The problem of strategic location choice for new antenna masts could be avoided if I was able to identify those antenna locations that network providers installed before the roll-out of 3G. Unfortunately, it is not possible to determine the construction date of every antenna location in Germany. I can therefore not reproduce the antenna infrastructure as it was before the 3G roll-out. However, every antenna location in Germany needs a certificate by the Federal Network Agency. The certification process ensures that all safety requirements with regards to electromagnetic fields are fulfilled and the potential exposure of persons living around the antenna site is contained (BEMFV, 2002). Each antenna site that fulfills the legal requirements is granted a certificate by the responsible branch of the Federal Network Agency. Until 2011, the Federal Network Agency had over 50 regional branches that were all responsible to issue certificates for antenna sites in their region. Every antenna site was given a sequential identification number which chronologically increased within each branch. Whenever new antennas are installed at an existing antenna site, the network provider has to apply for a new certificate but the identification number of the site does not change.

The sequential site identification numbers allow me to identify which sites are the oldest within each regional branch of the Federal Network Agency. My sample contains 63,481 antenna locations and I know that there were about 40,000 antenna locations before the roll-out of 3G. I could therefore conclude that the oldest 63 percent of antenna locations were likely to be in existence before the 3G upgrade. Of course, this is not a perfect way to determine the age of an antenna site since I only know the sequence within each of more than 50 regional branches of the Federal Network Agency. There might be some regions where more antennas had to be built and other regions where less antenna sites were built after the 3G upgrade. I therefore take a conservative approach and select the oldest 50 percent of antenna sites. It is highly likely that these 31,740 antenna locations were in operation before 3G and that these locations are not the result of strategic location choices made with the properties of 3G networks in mind. Figure 5.3 shows that the old antenna locations, depicted in green, are evenly distributed across Germany with a larger number of antenna sites in densely populated areas. The zoomed map section shows that even within densely populated areas, there is no obvious clustering of new antenna locations, depicted in red. I repeat the line-of-sight calculations with the restricted sample of old antenna masts.

Table 5.9 presents the estimations of Tables 5.2, 5.3 and 5.4 based on the reduced set of antenna locations. The first two columns show that the effect on the total number of reviews actually increases when visibility is calculated based on old antennas only. In the full sample, the coefficient is almost twice as large, in the rural sample it is even more than twice of the respective coefficient in

Table 5.2. Also, the probability that restaurants have at least one mobile review increases when only old antennas are taken into account. Compared to Table 5.3, the probability is 7 percentage points larger in the full sample and 3 percentage points larger in the rural sample. Only the effect on the share of mobile reviews over all reviews is slightly smaller when visibility to the restricted sample of antennas is used. While for the full sample of restaurants, I do not observe an effect for either visibility measure, in the rural sample the effect is reduced to 1.32 percentage points compared to 2.11 percentage points when the full set of antennas is used to calculate visibility.

### 5.6.2 Sensitivity to Different Sample Restrictions

For most analyses in this paper, I restrict my samples in different ways. In this section I show that neither of these restrictions drives the results. The arguably most arbitrary decision I am making concerns the definition of rural areas. As described in Section 5.2, I do not use standard administrative data about population density on the municipality or county level, but a more fine-grained definition based on a 1 by 1 kilometer population density grid. The reason for this choice is that the variation I use in 3G availability is rather within counties and municipalities than across counties and municipalities. There are few municipalities in Germany without any antenna masts and the average municipality has more than 5 masts. However, these masts are usually located in the center of a municipality. With the average municipality being about 29.5 square kilometers in size, there are many areas within a municipality that do not have 3G reception while there are few municipalities that do not have 3G reception as a whole. It is therefore not very practical to exclude whole municipalities based on their average population density.

On top of that, there are large differences in the size of municipalities across German states. For example, Hesse has an area of 21,115 square kilometers and 426 municipalities while Rhineland-Palatinate has an area of 19,853 square kilometers and 2,306 municipalities. Excluding municipalities based on their population density has therefore different implications in the two states. Since the variation I am using mostly occurs within municipalities and the administrative boundaries of municipalities are not consistent between states, it makes sense to use the fine-grained population density grid. The cutoff point for deciding if an area is rural or not is, however, subject for discussion. Generally, the density grid has a larger dispersion of density values than municipalities. For example, there are many areas with virtually zero population density, while, for obvious reasons, there is no municipality without population. At the other extreme, there are areas with a population density of



more than 15,000 inhabitants per square kilometer which is far beyond the most densely populated municipality.

In my sample, the average restaurant is located in an area with about 3,500 inhabitants per square kilometer. For all estimations, I arbitrarily called restaurants in areas of less than 2,500 inhabitants rural. In Table 5.10 I show how the effects on the total number of reviews and the share of mobile reviews over all reviews changes when I restrict the sample according to other definitions of rurality. Columns 2-4 show that compared to the original definition, the effect on total reviews does not become smaller in more restricted rural samples. Even in column 4, where I use only a third of the original rural observations and less than 8 percent of the full sample, my results stay robust. For the share of mobile reviews over all reviews, the effect becomes slightly smaller in the definitions of columns 6 and 7 but is the same in the most restricted definition in column 8. These results show that changing the definitions of rurality does not change the results in any meaningful way.

Another choice I am making is to restrict the sample to restaurants that had at least one review in 2008. The rationale behind this restriction is twofold. First, I want to compare restaurants that were in business before and after the treatment period. Since I define 2008 as the pre-treatment period, it makes sense to restrict the sample to restaurants that were in business 2008. Second, I want to compare restaurants that have customers with a similar propensity to use online services like Qype. By restricting the sample to restaurants that had reviews in the early days of the platform, I also hope to restrict the sample to restaurants with customers that are all principally open towards new technologies. However, one could also argue that the sample should be restricted to restaurants that had at least one review in 2007 instead of 2008.

Table 5.11 shows results for a restricted sample of restaurants that had at least one review in 2007. This restriction reduces the sample size to less than 11,000 restaurants or 45 percent of the sample of restaurants that had at least one review in 2008. In the rural sample, I am down to 2,362 restaurants or about 39 percent of the original rural sample. It can be assumed that the remaining restaurants have customers that are early adopters of online services like Qype. Albeit the large differences in sample size, I still find little differences in the effects on the total number of reviews, as shown in columns 1-2. In fact, the effects on total reviews are significantly larger compared to the effects in the baseline samples of Table 5.2. The effects on the probability of receiving at least one mobile review are virtually the same as the ones shown in Table 5.3. Only the effect on the

share of mobile over all reviews is about 20 percent smaller in the sample of restaurants that had at least one review in 2007.

In Section 5.4, I argue that 2008 is the appropriate base year for my analysis. The main reasons for this choice are that the iPhone 3G, arguably the first mass market 3G smartphone, was introduced in July 2008 together with the iOS App Store which allows to download third-party Apps. In late October 2008 the first Android smartphone launched together with Android Market, the App Store for Android devices. Figure 5.7 confirms that 2008 was also a turning point in the amount of data that was sent over the German cellular network. However, since the end of 2008 was already seeing growing sales of 3G smartphones and increasing data volumes, one could argue that it would be more appropriate to take 2007 as a base year for the difference-in-difference estimations.

Table 5.12 shows that the effects on the total number of reviews increase considerably when I take 2007 instead of 2008 as the pre-treatment period in my difference-in-difference estimations. The main reason is probably that in 2007 many restaurants that already had a few reviews in 2008 had zero reviews in 2007, as I can infer from the large sample reduction in Table 5.11 compared to Table 5.2. For the full sample, the coefficients in column 4 of Table 5.12 are more than twice as large as the respective coefficients in Table 5.2. For the rural sample, the change in the magnitude of the coefficients is even larger. For the outcomes that are related to mobile reviews, it is not necessary to repeat the estimations with the earlier base year as I observe the first mobile reviews in 2009, so that it does not make any difference if I compare 2012 with 2008 or 2007.

## 5.7 Conclusions

This paper provides evidence that upgrading the mobile Internet network from 2G to 3G increased the usage of local online services. Data from the review platform Qype show that restaurants in areas with a higher likelihood of 3G reception receive more reviews than restaurants in areas with a lower likelihood of 3G reception. By estimating a difference-in-difference model and exploiting technical differences between 2G and 3G signals in regard to their absorption by different kinds of terrain, the effect of the infrastructure upgrade can be interpreted as causal. In addition to the total number of reviews, also the probability of receiving at least one mobile review and the share of mobile reviews compared to non-mobile reviews has increased in areas with 3G. In rural areas, the effects are at least as large as in the whole sample, indicating that infrastructure upgrades can influence usage of the infrastructure even if the initial demand for the upgrade is low.

The effects of 3G availability on the number of restaurant reviews hold in the face of several robustness tests and sample restrictions. Given the reliance of our identification strategy on exogenous geographic variation and antenna locations, it would be worrisome if antenna locations were influenced to a large extent by the 3G roll-out itself. In order to rule out that my results are driven by the strategic choice of antenna locations, I restrict the antenna sample to locations that were largely existent before the 3G spectrum auction took place in Germany. The results do not only remain robust, but even become slightly larger for two out of three outcomes when only the oldest antenna sites are taken into account. I also show that the results are robust to a variety of sample restrictions that are related to my definition of rural areas, to the age of the restaurant and the first review they received, and to the base year for the difference-in-difference estimations.

This paper also provides first evidence that restaurant reviews which are submitted from mobile phones are not considerably more or less favorable than reviews which are submitted from PCs. Restaurants with a higher share of mobile reviews do therefore not have a better average rating. However, restaurants with a larger number of reviews might benefit from the increased consumer awareness and from increased consumer knowledge about their products. From a welfare perspective, a larger number of reviews is also associated with higher market transparency. The increase in transparency is in turn likely to promote competition and increase consumer welfare.

In order to justify subsidies for infrastructure upgrades in rural areas or other regions where market incentives are insufficient to attract private investments, it is necessary to analyze the returns of the infrastructure upgrade. Returns can only materialize if the upgrade of the infrastructure leads to a change of its usage. By showing that the upgrade from 2G to 3G leads to increased usage of local online services, this paper provides a first step in answering the chicken-and-egg question of infrastructure investments into mobile networks. It does not, however, show that the returns justify subsidies. Just for the application we are analyzing, it is probable that the increased usage of review platforms is associated with direct economic effects both for individual restaurants and for consumers more generally. The analysis of the economic returns to mobile network upgrades is subject to future research.

## Figures and Tables

Figure 5.1  
Screenshots from the Qype iPhone App

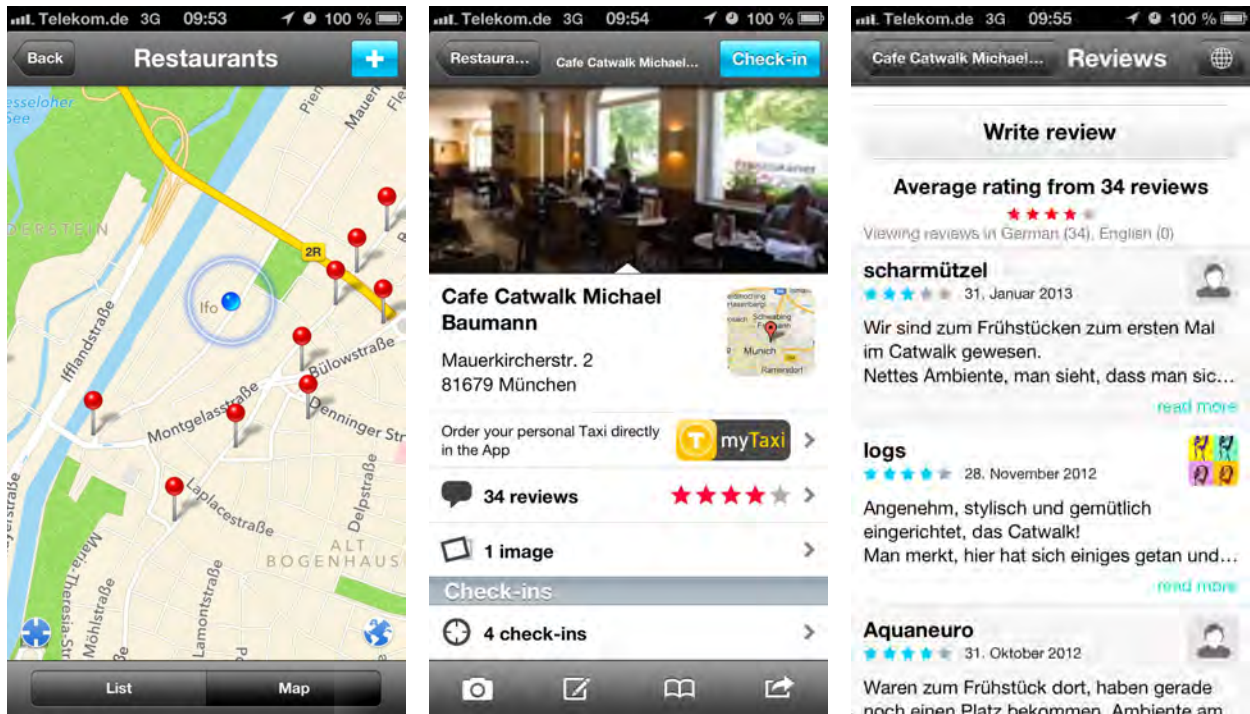


Figure 5.2  
All restaurants in Germany that can be found on Qype




Figure 5.3  
Old (green) and new (red) antenna locations in Germany




Figure 5.4  
Digital Surface Model (DSM) and antenna locations




Figure 5.5  
Disaggregated population density in the area around Munich




Figure 5.6  
Schematic representation of 2G and 3G penetration properties





Figure 5.7  
Data volumes in the German mobile network in million GB



Source: Annual Reports of the Bundesnetzagentur 2009-2011



Figure 5.8  
Number of reviews by visibility







Table 5.1  
Descriptive statistics in the full and in the rural sample

	Min	Max	Full sample			Rural areas		
			Mean	SD	Obs	Mean	SD	Obs
<b>Restaurant reviews</b>								
Total reviews in 2012	0	379	13.42	20.04	24054	8.02	12.14	6094
Total reviews in 2008	1	97	2.72	3.81	24192	1.92	2.25	6125
Mobile reviews	0	51	1.85	2.76	24054	1.27	2.24	6094
iPhone reviews	0	45	1.41	2.21	24192	0.97	1.79	6125
Android reviews	0	12	0.36	0.77	24192	0.24	0.61	6125
Share of mobile reviews	0	1	0.13	0.15	24049	0.13	0.17	6091
First review	2005	2008	2007.35	0.78	24192	2007.49	0.71	6125
Author experience	1	2782	93.54	149.56	24088	111.47	173.72	6072
<b>Antennas and geography</b>								
Visibility	0	1	0.75	0.43	24189	0.59	0.49	6122
Visible antennas	0	8	1.61	1.40	24192	1.09	1.23	6125
Antennas 10km	2	1978	422.18	485.68	24189	177.14	271.02	6122
Antennas 6km	1	1205	227.80	274.18	24189	83.96	145.05	6122
Distance closest	0	93	3.61	5.08	24192	7.53	8.11	6125
<b>Regional characteristics</b>								
Municipality population	0	3461	618.45	987.28	23624	217.73	564.67	5778
Municipality flaeche	0	1	0.24	0.27	23625	0.13	0.17	5779
Municipality density	6	4355	1615.65	1255.62	23624	829.51	989.60	5778
Density cell	0	15008	5056.04	3764.96	24192	636.82	901.03	6125
DSL coverage	0	99	90.98	8.70	23401	85.83	12.04	5919
<b>Restaurant characteristics</b>								
Claimed	0	1	0.24	0.43	24192	0.18	0.38	6125
Photo	0	1	0.67	0.47	24192	0.62	0.48	6125
Phone	0	1	0.96	0.20	24192	0.95	0.21	6125
Website	0	1	0.49	0.50	24192	0.49	0.50	6125
Online booking	0	1	0.06	0.23	24192	0.05	0.22	6125
Stars	0	5	3.85	0.83	24089	3.89	0.88	6101
Rank	1	2448	53.14	105.80	23462	24.14	60.42	6000
Checkins	0	546	2.81	9.15	24192	1.54	4.45	6125
German cuisine	0	1	0.20	0.40	24192	0.31	0.46	6125
Italian cuisine	0	1	0.17	0.38	24192	0.10	0.30	6125
Chinese cuisine	0	1	0.04	0.19	24192	0.03	0.17	6125
Greek cuisine	0	1	0.04	0.20	24192	0.03	0.17	6125

*Notes:* The table shows descriptive statistics of variables on the restaurant level in the full and rural samples. The full sample includes all restaurants that had at least one review in 2008. The rural sample includes restaurants in a 1-ha population density cell with less than 2500 inhabitants per square kilometer.

Table 5.2  
Effects on the total number of reviews

	Full sample				Rural areas			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Visibility × 2012	3.2662*** (0.2690)	3.0359*** (0.2276)	2.9741*** (0.2316)	3.2689*** (0.2547)	1.2144*** (0.3148)	1.2363*** (0.2624)	1.2099*** (0.2740)	1.2135*** (0.2719)
2012	8.2419*** (0.2169)	7.8028*** (0.1822)	8.1407*** (0.1858)	8.2360*** (0.2206)	5.3744*** (0.2325)	5.1121*** (0.1855)	5.3998*** (0.1947)	5.3707*** (0.2088)
Visibility	0.6387*** (0.0486)	-0.6083*** (0.0828)	-1.2251*** (0.0986)		0.3041*** (0.0551)	-0.1243 (0.0904)	-0.5714*** (0.1243)	
Restaurant characteristics	No	Yes	Yes	No	No	Yes	Yes	No
Restaurant types	No	Yes	Yes	No	No	Yes	Yes	No
Regional characteristics	No	No	Yes	No	No	No	Yes	No
County fixed effects	No	No	Yes	No	No	No	Yes	No
Restaurant fixed effects	No	No	No	Yes	No	No	No	Yes
Restaurants	24,120	23,245	21,954	24,120	6,107	5,919	5,404	6,107
Observations	48,240	46,489	43,908	48,240	12,213	11,837	10,808	12,213
R2	0.127	0.299	0.350	0.285	0.113	0.290	0.363	0.256

*Notes:* The dependent variable is the total number of reviews that a restaurant received until 2012. The full sample includes all restaurants with at least one review in 2008, the rural sample includes restaurants that are additionally located in a population density cell with less than 2500 inhabitants per square kilometer. Restaurant characteristics include controls like the star rating, the rank, the year of the first review, the experience of the average reviewer and the availability of photos. Restaurant types include more than 80 dummies that describe the cuisine or type of restaurant. Regional characteristics include controls like the number of antennas within 10km, the distance to the closest antenna, DSL availability and the population of the municipality. Ordinary least squares (OLS) estimations with robust standard errors clustered at the restaurant level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 5.3  
Effects on having at least one mobile review

	Full sample				Rural areas			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Visibility × 2012	0.0911*** (0.0074)	0.0876*** (0.0072)	0.0821*** (0.0073)	0.0911*** (0.0073)	0.0855*** (0.0129)	0.0842*** (0.0126)	0.0864*** (0.0131)	0.0855*** (0.0129)
2012	0.5235*** (0.0064)	0.5221*** (0.0063)	0.5334*** (0.0064)	0.5235*** (0.0063)	0.4367*** (0.0099)	0.4378*** (0.0097)	0.4484*** (0.0101)	0.4367*** (0.0099)
Visibility	-0.0000*** (0.0000)	-0.0177*** (0.0018)	-0.0296*** (0.0024)		0.0000*** (0.0000)	-0.0120*** (0.0031)	-0.0216*** (0.0049)	
Restaurant characteristics	No	Yes	Yes	No	No	Yes	Yes	No
Restaurant types	No	Yes	Yes	No	No	Yes	Yes	No
Regional characteristics	No	No	Yes	No	No	No	Yes	No
County fixed effects	No	No	Yes	No	No	No	Yes	No
Restaurant fixed effects	No	No	No	Yes	No	No	No	Yes
Restaurants	24,189	23,261	21,969	24,189	6,122	5,922	5,406	6,122
Observations	48,378	46,522	43,938	48,378	12,244	11,844	10,812	12,244
R2	0.424	0.478	0.498	0.595	0.327	0.389	0.433	0.491

Notes: The dependent variable is 1 if the restaurant has at least one mobile review by 2012. The full sample includes all restaurants with at least one review in 2008, the rural sample includes restaurants that are additionally located in a population density cell with less than 2500 inhabitants per square kilometer. Restaurant characteristics include controls like the star rating, the rank, the year of the first review, the experience of the average reviewer and the availability of photos. Restaurant types include more than 80 dummies that describe the cuisine or type of restaurant. Regional characteristics include controls like the number of antennas within 10km, the distance to the closest antenna, DSL availability and the population of the municipality. Ordinary least squares (OLS) estimations with robust standard errors clustered at the restaurant level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 5.4  
Effects on the share of mobile reviews out of all reviews

	Full sample				Rural areas			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Visibility × 2012	0.0080*** (0.0024)	0.0077*** (0.0024)	0.0066*** (0.0024)	0.0080*** (0.0023)	0.0211*** (0.0044)	0.0210*** (0.0044)	0.0217*** (0.0045)	0.0211*** (0.0044)
2012	0.1273*** (0.0021)	0.1280*** (0.0021)	0.1289*** (0.0021)	0.1273*** (0.0020)	0.1176*** (0.0033)	0.1187*** (0.0033)	0.1196*** (0.0034)	0.1176*** (0.0034)
Visibility	0.0000*** (0.0000)	-0.0008** (0.0003)	-0.0013** (0.0006)		0.0000 (0.0000)	-0.0027*** (0.0008)	-0.0039*** (0.0015)	
Restaurant characteristics	No	Yes	Yes	No	No	Yes	Yes	No
Restaurant types	No	Yes	Yes	No	No	Yes	Yes	No
Regional characteristics	No	No	Yes	No	No	No	Yes	No
County fixed effects	No	No	Yes	No	No	No	Yes	No
Restaurant fixed effects	No	No	No	Yes	No	No	No	Yes
Restaurants	24,118	23,243	21,953	24,118	6,105	5,918	5,404	6,105
Observations	48,235	46,486	43,906	48,235	12,210	11,836	10,807	12,210
R2	0.275	0.292	0.314	0.431	0.234	0.264	0.311	0.377

*Notes:* The dependent variable is the share of mobile reviews out of all reviews until 2012. The full sample includes all restaurants with at least one review in 2008, the rural sample includes restaurants that are additionally located in a population density cell with less than 2500 inhabitants per square kilometer. Restaurant characteristics include controls like the star rating, the rank, the year of the first review, the experience of the average reviewer and the availability of photos. Restaurant types include more than 80 dummies that describe the cuisine or type of restaurant. Regional characteristics include controls like the number of antennas within 10km, the distance to the closest antenna, DSL availability and the population of the municipality. Ordinary least squares (OLS) estimations with robust standard errors clustered at the restaurant level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 5.5  
The source of reviews and their corresponding star rating

	(1)	(2)	(3)	(4)	(5)	(6)
From mobile phone	0.1100*** (0.0053)	0.1009*** (0.0052)	0.1215*** (0.0052)	0.1051*** (0.0055)	0.1061*** (0.0054)	0.0998*** (0.0054)
Comments		-0.1166*** (0.0017)	-0.1189*** (0.0017)	-0.1195*** (0.0018)		-0.1219*** (0.0017)
Compliments		0.0101*** (0.0003)	0.0096*** (0.0003)	0.0098*** (0.0003)		0.0089*** (0.0003)
Author experience		-0.0006*** (0.0000)	-0.0006*** (0.0000)	-0.0005*** (0.0001)		-0.0003*** (0.0001)
Years	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant characteristics	No	No	Yes	Yes	No	No
Restaurant types	No	No	Yes	Yes	No	No
Regional characteristics	No	No	No	Yes	No	No
County fixed effects	No	No	No	Yes	No	No
Restaurant fixed effects	No	No	No	No	Yes	Yes
Restaurants	65,848	65,848	65,848	59,798	65,848	65,848
Observations	459,748	459,748	459,748	428,995	459,748	459,748
R2	0.004	0.014	0.031	0.038	0.011	0.023

Notes: The dependent variable is the star rating of the review from 1 to 5. Restaurant characteristics include controls like the star rating, the rank, the year of the first review, the experience of the average reviewer and the availability of photos. Restaurant types include more than 80 dummies that describe the cuisine or type of restaurant. Regional characteristics include controls like the number of antennas within 10km, the distance to the closest antenna, DSL availability and the population of the municipality. Ordinary least squares (OLS) estimations with robust standard errors clustered at the restaurant level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 5.6  
The source and star rating of reviews with restaurant and author fixed effects

	Within restaurants			Within authors		
	(1)	(2)	(3)	(4)	(5)	(6)
From mobile phone	0.0185* (0.0099)	0.0213** (0.0099)	-0.0646*** (0.0101)	-0.0652*** (0.0101)	-0.0594*** (0.0101)	-0.0556*** (0.0100)
Comments		-0.0394*** (0.0025)		-0.0393*** (0.0024)	-0.0398*** (0.0024)	-0.0398*** (0.0024)
Compliments		0.0046*** (0.0005)		0.0074*** (0.0007)	0.0067*** (0.0007)	0.0062*** (0.0007)
Author experience		0.0001* (0.0001)				
Years	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant characteristics	No	No	No	No	Yes	Yes
Restaurant types	No	No	No	No	No	Yes
Restaurants	35,240	35,240	35,240	35,240	35,240	35,240
Authors	8,660	8,660	8,660	8,660	8,660	8,660
Observations	93,746	93,746	93,746	93,746	93,746	93,746
R2	0.00	0.01	0.00	0.00	0.01	0.03

Notes: The dependent variable is the star rating of the review from 1 to 5. Column (1)-(2) include restaurant-fixed effects, columns (3)-(6) include author-fixed effects. For comparability, the samples are reduced to reviews from authors that gave at least one mobile and at least one non-mobile review. Restaurant characteristics include controls like the star rating, the rank, the year of the first review, the experience of the average reviewer and the availability of photos. Restaurant types include more than 80 dummies that describe the cuisine or type of restaurant. Ordinary least squares (OLS) estimations with robust standard errors clustered at the restaurant level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 5.7  
Mobile phone brands and the star rating of reviews

	Within restaurants				Within authors		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
From iPhone	0.0703*** (0.0059)			0.0815*** (0.0059)	-0.0605*** (0.0115)		
From Android		0.1279*** (0.0108)		0.1468*** (0.0109)		-0.0328* (0.0196)	
From BlackBerry			0.1669*** (0.0246)	0.1923*** (0.0246)			-0.0128 (0.0479)
Comments	-0.1222*** (0.0017)	-0.1223*** (0.0017)	-0.1224*** (0.0017)	-0.1219*** (0.0017)	-0.0793*** (0.0038)	-0.0170*** (0.0032)	-0.1006*** (0.0220)
Compliments	0.0089*** (0.0003)	0.0089*** (0.0003)	0.0089*** (0.0003)	0.0089*** (0.0003)	0.0067*** (0.0009)	0.0062*** (0.0014)	0.0036 (0.0034)
Author experience	-0.0003*** (0.0001)	-0.0003*** (0.0001)	-0.0003*** (0.0001)	-0.0003*** (0.0001)			
Years	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant characteristics	No	No	No	No	Yes	Yes	Yes
Restaurant types	No	No	No	No	Yes	Yes	Yes
Restaurants	65,848	65,848	65,848	65,848	30,376	15,840	3,529
Authors	151,145	151,145	151,145	151,145	6,476	2,434	346
Observations	459,748	459,748	459,748	459,748	71,269	25,228	3,910
R2	0.02	0.02	0.02	0.02	0.03	0.03	0.06

*Notes:* The dependent variable is the star rating of the review from 1 to 5. Column (1)-(4) include restaurant-fixed effects, columns (5)-(7) include author-fixed effects. In columns (5)-(7), the samples are reduced to reviews from authors that gave at least one review from the respective mobile phone brand and at least one non-mobile review. Restaurant characteristics include controls like the star rating, the rank, the year of the first review, the experience of the average reviewer and the availability of photos. Restaurant types include more than 80 dummies that describe the cuisine or type of restaurant. Ordinary least squares (OLS) estimations with robust standard errors clustered at the restaurant level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.



Table 5.8  
The number of mobile and non-mobile reviews and the overall star rating of a restaurant

	Full sample				Rural areas			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
No. of mobile reviews	-0.0038* (0.0020)	-0.0045** (0.0020)	-0.0055*** (0.0020)	-0.0022* (0.0013)	-0.0051 (0.0050)	-0.0023 (0.0049)	-0.0041 (0.0050)	-0.0031 (0.0030)
No. of PC reviews	-0.0022*** (0.0003)	-0.0042*** (0.0003)	-0.0020*** (0.0003)	-0.0036*** (0.0003)	-0.0077*** (0.0010)	-0.0119*** (0.0011)	-0.0071*** (0.0012)	-0.0088*** (0.0010)
Years	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant characteristics	No	Yes	Yes	No	No	Yes	Yes	No
Restaurant types	No	Yes	Yes	No	No	Yes	Yes	No
Regional characteristics	No	No	Yes	No	No	No	Yes	No
County fixed effects	No	No	Yes	No	No	No	Yes	No
Restaurant fixed effects	No	No	No	Yes	No	No	No	Yes
Restaurants	24,114	24,012	22,684	24,114	6,102	6,051	5,522	6,102
Observations	120,568	120,058	113,418	120,568	30,510	30,255	27,610	30,510
R2	0.007	0.042	0.083	0.043	0.008	0.051	0.143	0.041

*Notes:* The dependent variable is the average star rating of a restaurant, ranging from 1 to 5. The full sample includes all restaurants with at least one review in 2008, the rural sample includes restaurants that are additionally located in a population density cell with less than 2500 inhabitants per square kilometer. Restaurant characteristics include controls like the star rating, the rank, the year of the first review, the experience of the average reviewer and the availability of photos. Restaurant types include more than 80 dummies that describe the cuisine or type of restaurant. Regional characteristics include controls like the number of antennas within 10km, the distance to the closest antenna, DSL availability and the population of the municipality. Ordinary least squares (OLS) estimations with robust standard errors clustered at the restaurant level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 5.9  
Effects when only the oldest 50 percent of antennas are used to determine visibility

	Total reviews		Any mobile review		Mobile reviews share	
	Full (1)	Rural (2)	Full (3)	Rural (4)	Full (5)	Rural (6)
Visibility × 2012	6.3975*** (0.1980)	2.7133*** (0.2592)	0.1621*** (0.0085)	0.1161*** (0.0133)	0.0074** (0.0029)	0.0132*** (0.0046)
2012	5.0655*** (0.1464)	4.4290*** (0.1655)	0.4607*** (0.0078)	0.4283*** (0.0105)	0.1234*** (0.0027)	0.1190*** (0.0037)
Visibility	-3.2393*** (0.1152)	-1.3399*** (0.1431)	-0.0687*** (0.0036)	-0.0435*** (0.0059)	-0.0010 (0.0011)	-0.0061*** (0.0018)
Restaurant characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant types	Yes	Yes	Yes	Yes	Yes	Yes
Regional characteristics	Yes	Yes	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Restaurants	21,905	5,358	21,920	5,360	21,904	5,358
Observations	43,810	10,716	43,840	10,720	43,808	10,715
R2	0.354	0.366	0.501	0.437	0.309	0.305

Notes: The dependent variable in columns (1)-(2) is the total number of reviews, in columns (3)-(4) it is having at least mobile review and in columns (5)-(6) it is the share of mobile reviews out of all reviews until 2012. Only the oldest 50 percent of antennas are used for visibility calculations. Restaurant characteristics include controls such as the star rating, the rank, the year of the first review, the experience of the average reviewer and the availability of photos. Regional characteristics include controls such as the number of antennas within 10km, the distance to the closest antenna, DSL availability and the population of the municipality. Ordinary least squares (OLS) estimations with robust standard errors clustered at the restaurant level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 5.10  
Effects in samples with different rural area definitions

	Total number of reviews				Share of mobile reviews			
	< 2500 (1)	< 1500 (2)	< 500 (3)	< 100 (4)	< 2500 (5)	< 1500 (6)	< 500 (7)	< 100 (8)
Visibility × 2012	1.2099*** (0.2700)	1.6289*** (0.3423)	1.7087*** (0.3587)	1.3886*** (0.5321)	0.0217*** (0.0047)	0.0162*** (0.0054)	0.0187*** (0.0055)	0.0219*** (0.0060)
2012	5.3998*** (0.1942)	5.9115*** (0.2408)	5.8791*** (0.2557)	6.1241*** (0.3764)	0.1196*** (0.0036)	0.1217*** (0.0041)	0.1190*** (0.0043)	0.1172*** (0.0046)
Visibility	-0.5714*** (0.1222)	-0.8359*** (0.1614)	-0.8378*** (0.1715)	-0.9215*** (0.2736)	-0.0039*** (0.0015)	-0.0048*** (0.0018)	-0.0056*** (0.0019)	-0.0069*** (0.0021)
Restaurant characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant types	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Regional characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Restaurants	5,404	4,084	3,819	1,856	5,404	4,084	3,819	3,198
Observations	10,808	8,168	7,638	3,711	10,807	8,168	7,638	6,396
R2	0.363	0.373	0.373	0.413	0.311	0.323	0.325	0.331

Notes: The dependent variable in columns (1)-(4) is the total number of reviews and in columns (5)-(8) it is the share of mobile reviews out of all reviews until 2012. The samples are restricted to restaurants that are located in a population density cell with less than 2500, 1500, 500 and 100 inhabitants per square kilometer, respectively. Restaurant characteristics include controls like the star rating, the rank, the year of the first review, the experience of the average reviewer and the availability of photos. Restaurant types include more than 80 dummies that describe the cuisine or type of restaurant. Regional characteristics include controls like the number of antennas within 10km, the distance to the closest antenna, DSL availability and the population of the municipality. Ordinary least squares (OLS) estimations with robust standard errors clustered at the restaurant level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 5.11  
Effects for a restricted sample of restaurants that had at least one review in 2007

	Total reviews		Any mobile review		Mobile reviews share	
	Full (1)	Rural (2)	Full (3)	Rural (4)	Full (5)	Rural (6)
visitoday	3.9511*** (0.4545)	2.1220*** (0.5795)	0.0865*** (0.0108)	0.0906*** (0.0212)	0.0023 (0.0033)	0.0170** (0.0068)
today	11.5134*** (0.3750)	7.1297*** (0.4031)	0.5984*** (0.0096)	0.4922*** (0.0166)	0.1259*** (0.0030)	0.1191*** (0.0053)
visi	-1.8100*** (0.1953)	-0.9269*** (0.2941)	-0.0350*** (0.0039)	-0.0243** (0.0095)	-0.0004 (0.0009)	-0.0033 (0.0027)
Restaurant characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant types	Yes	Yes	Yes	Yes	Yes	Yes
Regional characteristics	Yes	Yes	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Restaurants	9,950	2,077	9,953	2,077	9,949	2,077
Observations	19,899	4,154	19,906	4,154	19,898	4,154
R2	0.372	0.425	0.573	0.499	0.359	0.366

Notes: The dependent variable in columns (1)-(2) is the total number of reviews, in columns (3)-(4) it is having at least mobile review and in columns (5)-(6) it is the share of mobile reviews out of all reviews until 2012. The full sample includes all restaurants with at least one review in 2007, the rural sample includes restaurants that are additionally located in a population density cell with less than 2500 inhabitants per square kilometer. Restaurant characteristics include controls such as the star rating, the rank, the year of the first review, the experience of the average reviewer and the availability of photos. Regional characteristics include controls such as the number of antennas within 10km, the distance to the closest antenna, DSL availability and the population of the municipality. Ordinary least squares (OLS) estimations with robust standard errors clustered at the restaurant level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.

Table 5.12  
Effects on the total number of reviews when 2007 is used as the base year in difference-in-difference estimations

	Full sample				Rural areas			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Visibility × 2012	7.7142*** (0.2132)	7.1083*** (0.1939)	7.1349*** (0.1977)	7.7142*** (0.3209)	3.3725*** (0.2828)	3.0712*** (0.2478)	3.1201*** (0.2604)	3.3739*** (0.3033)
2012	6.1599*** (0.1494)	6.0083*** (0.1422)	6.2685*** (0.1454)	6.1581*** (0.2919)	5.3654*** (0.1656)	5.2865*** (0.1558)	5.5346*** (0.1658)	5.3632*** (0.2389)
Visibility	0.4506*** (0.0193)	-1.9639*** (0.0929)	-3.6215*** (0.1139)		0.1712*** (0.0269)	-0.5866*** (0.0883)	-1.5691*** (0.1424)	
Restaurant characteristics	No	Yes	Yes	No	No	Yes	Yes	No
Restaurant types	No	Yes	Yes	No	No	Yes	Yes	No
Regional characteristics	No	No	Yes	No	No	No	Yes	No
County fixed effects	No	No	Yes	No	No	No	Yes	No
Restaurant fixed effects	No	No	No	Yes	No	No	No	Yes
Restaurants	24,059	23,188	21,905	24,059	6,049	5,865	5,358	6,049
Observations	48,118	46,375	43,810	48,118	12,097	11,729	10,716	12,097
R2	0.183	0.322	0.363	0.320	0.175	0.323	0.380	0.308

*Notes:* The dependent variable is the total number of reviews that a restaurant received until 2012. The pre-treatment period is 2007. The full sample includes all restaurants with at least one review in 2008, the rural sample includes restaurants that are additionally located in a population density cell with less than 2500 inhabitants per square kilometer. Restaurant characteristics include controls like the star rating, the rank, the year of the first review, the experience of the average reviewer and the availability of photos. Restaurant types include more than 80 dummies that describe the cuisine or type of restaurant. Regional characteristics include controls like the number of antennas within 10km, the distance to the closest antenna, DSL availability and the population of the municipality. Ordinary least squares (OLS) estimations with robust standard errors clustered at the restaurant level in parentheses. Significance levels denoted by \* at 10%, \*\* at 5% and \*\*\* at 1%.



---

## References

---

- Addison, John T., and Pedro Portugal.** 2001. "Job Search Methods and Outcomes." *IZA Discussion Paper*, 349.
- Ahlfeldt, Gabriel, Pantelis Koutroumpis, and Tommaso Valletti.** 2014. "Speed 2.0 - Evaluating Access to Universal Digital Highways." *mimeo*.
- Ai, Chunrong, and Edward C. Norton.** 2003. "Interaction terms in logit and probit models." *Economics Letters*, 80(1): 123–129.
- Altinok, Nadir, and Geeta Kingdon.** 2012. "New Evidence on Class Size Effects: A Pupil Fixed Effects Approach." *Oxford Bulletin of Economics and Statistics*, 74(2): 203–234.
- Altmann, Jörn, Björn Rupp, and Pravin Varaiya.** 2001. "Effects of pricing on Internet user behavior." *Netnomics*, 3(1): 67–84.
- Ammermüller, Andreas, and Peter Dolton.** 2006. "Pupil-Teacher Gender Interaction Effects on Scholastic Outcomes in England and the USA." *ZEW Discussion Paper*, 06-60.
- Anderson, Michael, and Jeremy Magruder.** 2012. "Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database." *Economic Journal*, 122(563): 957–989.
- Angrist, J, and JS Pischke.** 2010. "The credibility revolution in empirical economics: How better research design is taking the con out of econometrics." *Journal of Economic Perspectives*, 24(2): 3–30.
- Angrist, Joshua D., and Jörn-Steffen Pischke.** 2008. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ:Princeton University Press.
- Angrist, Joshua D., and Victor Lavy.** 2002. "New Evidence on Classroom Computers and Pupil Learning." *Economic Journal*, 112(482): 735–765.
- Aschauer, David A.** 1989. "Is public expenditure productive?" *Journal of Monetary Economics*, 23(2): 177–200.
- Ashenfelter, Orley C., and David J. Zimmerman.** 1997. "Estimates of the returns to schooling from sibling data: Fathers, sons, and brothers." *Review of Economics and Statistics*, 79(1): 1–9.
- Aslam, Monazza, and Geeta Kingdon.** 2011. "What can teachers do to raise pupil achievement?" *Economics of Education Review*, 30(3): 559–574.
- Atack, Jeremy, Fred Bateman, M. Haines, and Robert A. Margo.** 2010. "Did railroads induce or follow economic growth? Urbanization and Population Growth in the American Midwest, 1850-1860." *Social Science History*, 34(2): 171–197.

- Attewell, Paul, and Juan Battle.** 1999. "Home Computers and School Performance." *The Information Society*, 15(1): 1–10.
- Autor, David H.** 2001. "Wiring the Labor Market." *Journal of Economic Perspectives*, 15(1): 25–40.
- Bajari, Patrick, and Ali Hortacsu.** 2004. "Economic Insights from Internet Auctions." *Journal of Economic Literature*, 42(2): 457–486.
- Banerjee, Abhijit, Cole Shawn, Esther Duflo, and Leigh Linden.** 2007. "Remedying education: evidence from two randomized experiments in India." *Quarterly Journal of Economics*, 1235–1265.
- Banerjee, Abhijit, Esther Duflo, and Nancy Qian.** 2012. "On the Road: Access to Transportation Infrastructure and Economic Growth in China." *NBER Working Paper*, 17897.
- Barrera-Osorio, Felipe, and Leigh Linden.** 2009. "The Use and Misuse of Computers in Education: Evidence from a Randomized Experiment in Colombia." *Policy Research Working Paper*, 4836.
- Bauernschuster, Stefan, Oliver Falck, and Ludger Woessmann.** 2014. "Surfing Alone? The Internet and Social Capital - Evidence from an Unforeseeable Technological Mistake." *Journal of Public Economics*, forthcoming.
- Bauer, Thomas K., Sven Feuerschütte, Michael Kiefer, Philipp an de Meulen, Martin Micheli, Torsten Schmidt, and Lars-Holger Wilke.** 2013. "Ein hedonischer Immobilienpreisindex auf Basis von Internetdaten: 2007-2011." *AStA Wirtschafts-und Sozialstatistisches Archiv*, 7(1-2): 5–30.
- Belo, Rodrigo, Pedro Ferreira, and Rahul Telang.** 2013. "Broadband in school: Impact on student performance." *Management Science*, 60(2): 265–282.
- Beltran, Daniel O., Kuntal K. Das, and Robert W. Fairlie.** 2010. "Home Computers and Educational Outcomes: Evidence from the NLSY97 and CPS." *Economic Inquiry*, 48(3): 771–792.
- BEMFV.** 2002. *Verordnung über das Nachweisverfahren zur Begrenzung elektromagnetischer Felder vom 20. August 2002 (BGBl. I S. 3366).*
- Berger, Jonah, Alan T. Sorensen, and Scott J. Rasmussen.** 2010. "Positive Effects of Negative Publicity: When Negative Reviews Increase Sales." *Marketing Science*, 29(5): 815–827.
- Black, Sandra E.** 1999. "Do Better Schools Matter? Parental Valuation of Elementary Education." *Quarterly Journal of Economics*, 114(2): 577–599.
- Bowlus, Audra J.** 1995. "Matching Workers and Jobs: Cyclical Fluctuations in Match Quality." *Journal of Labor Economics*, 13(2): 335–350.
- Bresnahan, Timothy F., and Manuel Trajtenberg.** 1995. "General purpose technologies 'Engines of growth'?" *Journal of Econometrics*, 65(1): 83–108.
- Bresnahan, Timothy F., and Robert J. Gordon.** 1996. "Introduction to "The Economics of New Goods"." In *The Economics of New Goods*. Vol. I, , ed. Timothy F. Bresnahan and Robert J. Gordon, 1–26. Chicago:University of Chicago Press.
- Brynjolfsson, Erik, and Adam Saunders.** 2009. "What the GDP Gets Wrong (Why Managers Should Care)." *MIT Sloan Management Review*, 51(1): 95–96.



- Brynjolfsson, Erik, and Joo H. Oh.** 2012. "The Attention Economy: Measuring the Value of Free Digital Services on the Internet." *mimeo*.
- Brynjolfsson, Erik, Yu Hu, and Michael D. Smith.** 2003. "Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers." *Management Science*, 49(11): 1580–1596.
- Bughin, Jacques.** 2011. "The Web's EUR 100 billion surplus." *McKinsey Quarterly*.
- Bughin, Jacques, and James Manyika.** 2014. "The mobile Internet's consumer dividend." *McKinsey Quarterly*.
- Bui, Linda T. M., and Christopher J. Mayer.** 2003. "Regulation and Capitalization of Environmental Amenities: Evidence from the Toxic Release Inventory in Massachusetts." *The Review of Economics and Statistics*, 85(3): 693–708.
- Bundesnetzagentur.** 2008. "Jahresbericht 2008."
- Bundesnetzagentur.** 2010. "Frequenznutzungsbestimmungen LTE 800 Mhz."
- Cabral, L., and A. Hortacsu.** 2010. "The Dynamics of Seller Reputation: Evidence from eBay." *The Journal of Industrial Economics*, 58(1): 54–79.
- Cai, Hongbin, Yuyu Chen, and Hanming Fang.** 2009. "Observational learning: Evidence from a randomized natural field experiment." *American Economic Review*, 99(3): 864–882.
- Cairncross, Frances.** 1997. *The Death of Distance: How the Communications Revolution Will Change Our Lives*. Boston, USA:Harvard Business School Press.
- Centeno, Mario.** 2004. "The Match Quality Gains from Unemployment Insurance." *Journal of Human Resources*, 39(3): 839–863.
- Chamberlain, Gary.** 1982. "Multivariate regression models for panel data." *Journal of Econometrics*, 18(1): 5–46.
- Chay, Kenneth Y., and Michael Greenstone.** 2005. "Does Air Quality Matter? Evidence from the Housing Market Michael Greenstone." *Journal of Political Economy*, 113(2): 376–424.
- Chen, Yan, Grace Y. J. Jeon, and Yong-Mi Kim.** 2014. "A day without a search engine: an experimental study of online and offline searches." *Experimental Economics*, forthcoming.
- Chevalier, Judith A., and Dina Mayzlin.** 2006. "The effect of word of mouth on sales: Online book reviews." *Journal of Marketing Research*, 43(3): 345–354.
- Chintagunta, Pradeep K., Shyam Gopinath, and Sriram Venkataraman.** 2010. "The Effects of Online User Reviews on Movie Box Office Performance: Accounting for Sequential Rollout and Aggregation Across Local Markets." *Marketing Science*, 29(5): 944–957.
- Clark, William.** 1988. "Production costs and output qualities in public and private employment agencies." *Journal of Law and Economics*, 31(2): 379–393.
- Clotfelter, Charles T., Helen F. Ladd, and Jacob L. Vigdor.** 2010. "Teacher credentials and student achievement in high school a cross-subject analysis with student fixed effects." *Journal of Human Resources*, 45(3): 655–681.
- Cristia, Julián, Ana Santiago, Eugenio Severín, and Pablo Ibararán.** 2012. "Technology and Child Development : Evidence from the One Laptop per Child Program." *IZA Discussion Papers*, 6401.

- Czernich, Nina.** 2012. "Broadband Internet and Political Participation: Evidence for Germany." *KYKLOS*, 65(1): 31–52.
- Czernich, Nina, Oliver Falck, Tobias Kretschmer, and Ludger Woessmann.** 2011. "Broadband Infrastructure and Economic Growth." *Economic Journal*, 121(552): 505–532.
- Dai, Weijia, Ginger Z. Jin, Jungmin Lee, and Michael Luca.** 2012. "Optimal Aggregation of Consumer Ratings: An Application to Yelp. com." *NBER Working Paper*, 18567.
- Davis, Lucas W.** 2004. "The Effect of Health Risk on Housing Values : Evidence from a Cancer Cluster." *American Economic Review*, 94(5): 1693–1704.
- Dee, Thomas S.** 2005. "A teacher like me: Does race, ethnicity, or gender matter?" *American Economic Review*, 95(2): 158–165.
- Dee, Thomas S.** 2007. "Teachers and the gender gaps in student achievement." *Journal of Human Resources*, 42(3): 528–554.
- Dee, Thomas S., and Martin R. West.** 2011. "The Non-Cognitive Returns to Class Size." *Educational Evaluation and Policy Analysis*, 33(1): 23–46.
- Dehejia, Rajeev H., and Sadek Wahba.** 2002. "Propensity score-matching methods for nonexperimental causal studies." *Review of Economics and Statistics*, 84(1): 151–161.
- Diamond, Peter A., and Jerry A. Hausman.** 1994. "Contingent Valuation: Is Some Number Better than No Number?" *Journal of Economic Perspectives*, 8(4): 45–64.
- Duranton, Gilles, and Matthew A. Turner.** 2011. "The fundamental law of road congestion: Evidence from US cities." *American Economic Review*, 101(6): 2616–2652.
- Dutz, Mark, Jonathan Orszag, and Robert Willig.** 2009. "The substantial consumer benefits of broadband connectivity for US households." *Internet Innovation Alliance Working Paper*.
- Ellison, Glenn, and Sara F. Ellison.** 2009. "Search, obfuscation, and price elasticities on the internet." *Econometrica*, 77(2): 427–452.
- Fairlie, Robert W.** 2004. "Race and the Digital Divide." *Berkeley Electronic Journals, Contributions to Economic Analysis and Policy*, 3(1).
- Fairlie, Robert W.** 2005. "The effects of home computers on school enrollment." *Economics of Education Review*, 24(5): 533–547.
- Fairlie, Robert W., and Jonathan Robinson.** 2013. "Experimental Evidence on the Effects of Home Computers on Academic Achievement among Schoolchildren." *American Economic Journal: Applied Economics*, 5(3): 211–240.
- Falch, Torberg, and Marte Ronning.** 2012. "Homework assignment and student achievement in OECD countries." *Statistics Norway Discussion Papers*, 711.
- Falck, Oliver, Robert Gold, and Stephan Heblich.** 2014. "E-lections: Voting Behavior and the Internet." *American Economic Review*, forthcoming.
- Farr, Tom G., Paul A. Rosen, Edward Caro, Robert Crippen, Riley Duren, Scott Hensley, Michael Kobrick, Mimi Paller, Ernesto Rodriguez, Ladislav Roth, David Seal, Scott Shaffer, Joanne Shimada, Jeffrey Umland, Marian Werner, Michael Oskin, Douglas Burbank, and Douglas Alsdorf.** 2007. "The Shuttle Radar Topography Mission." *Reviews of Geophysics*, 45(2): 1–33.

- FCC.** 2011. *Report and Order and Further Notice of Proposed Rulemaking 11-161*.
- Fernald, John G.** 1999. "Roads to Prosperity ? Assessing the Link Between Public Capital and Productivity." *American Economic Review*, 89(3): 619–38.
- Ferreira, Priscila, and Mark Taylor.** 2011. "Measuring match quality using subjective data." *Economics Letters*, 113(3): 304–306.
- Figlio, David N., and Maurice E. Lucas.** 2004. "What's in a Grade? School Report Cards and the Housing Market." *American Economic Review*, 94(3): 591–604.
- Fiorini, Mario.** 2010. "The effect of home computer use on childrens cognitive and non-cognitive skills." *Economics of Education Review*, 29(1): 55–72.
- Fogel, Robert W.** 1962. "A Quantitative Approach to the Study of Railroads in American Economic Growth: A Report of Some Preliminary Findings." *Journal of Economic History*, 22(2): 163–197.
- Forman, Chris, Avi Goldfarb, and Shane Greenstein.** 2012. "The Internet and Local Wages: A Puzzle." *American Economic Review*, 102(1): 556–575.
- Freeman, Richard B.** 2002. "The labour market in the new information economy." *Oxford Review of Economic Policy*, 18(3): 288–305.
- Fuchs, Thomas, and Ludger Woessmann.** 2004. "Computers and Student Learning: Bivariate and Multivariate Evidence on the Availability and use of Computers at Home and at School." *Brussels Economic Review*, 47(3/4): 359–385.
- Gallego, Francisco J.** 2010. "A population density grid of the European Union." *Population and Environment*, 31(6): 460–473.
- Gautier, Pieter A., and Christian L. Holzner.** 2011. "Simultaneous search and network efficiency." *IZA Discussion Papers*, 5859.
- Gayer, Ted, James T Hamilton, and W. Kip Viscusi.** 2000. "Private Values of Risk Tradeoffs at Superfund Sites: Housing Market Evidence Learning about Risk." *Review of Economics and Statistics*, 82(3): 439–451.
- Gebhardt, Georg.** 2010. "The Impact of the Internet on Retail Competition: Evidence from Technological Differences in Internet Access." *mimeo*.
- Georgi, Sabine, and Peter Barkow.** 2010. *Wohnimmobilien-Indizes : Vergleich Deutschland - Grossbritannien*. Berlin:Zentraler Immobilien Ausschuss e.V., ZIA.
- Gibbons, Steve, and Stephen Machin.** 2001. "Valuing Primary Schools."
- Goldfarb, Avi, and Jeff Prince.** 2008. "Internet adoption and usage patterns are different: Implications for the digital divide." *Information Economics and Policy*, 20(1): 2–15.
- Google Inc.** 2012. "Google s Economic Impact."
- Goolsbee, Austan.** 2006. "The value of broadband and the deadweight loss of taxing new technology." *B.E. Journal of Economic Analysis & Policy*, 5(1): 1–29.
- Goolsbee, Austan, and Jeffrey R. Brown.** 2002. "Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry." *Journal of Political Economy*, 110(3): 481–507.

- Goolsbee, Austan, and Jonathan Guryan.** 2006. "The Impact of Internet Subsidies in Public Schools." *Review of Economics and Statistics*, 88(2): 336–347.
- Goolsbee, Austan, and P.J. Klenow.** 2006. "Valuing consumer products by the time spent using them: an application to the internet." *American Economic Review*, 2(96): 108–113.
- Greenstein, Shane, and Ryan C. McDevitt.** 2011a. "Evidence of a modest price decline in US broadband services." *Information Economics and Policy*, 23(2): 200–211.
- Greenstein, Shane, and Ryan C. McDevitt.** 2011b. "The broadband bonus: Estimating broadband Internet's economic value." *Telecommunications Policy*, 35(7): 617–632.
- Greenstein, Shane, and Ryan C. McDevitt.** 2012. "Measuring the Broadband Bonus in Thirty OECD Countries." *OECD Digital Economy Papers*, 197.
- Greenstone, Michael, and Justin Gallagher.** 2008. "Does hazardous waste matter? Evidence from the housing market and the superfund program." *Quarterly Journal of Economics*, 123(3): 951–1003.
- Hallstrom, Daniel G., and V. Kerry Smith.** 2005. "Market responses to hurricanes." *Journal of Environmental Economics and Management*, 50(3): 541–561.
- Hanushek, Eric A.** 1997. "Assessing the effects of school resources on student performance: An update." *Educational Evaluation and Policy Analysis*, 19(2): 141–164.
- Hanushek, Eric A.** 2003. "The Failure of Input-based Schooling Policies." *Economic Journal*, 113(485): F64–F98.
- Hanushek, Eric A., and Ludger Woessmann.** 2011. "The Economics of International Differences in Educational Achievement." In *Handbook of the Economics of Education*. Vol. 3. 1 ed., , ed. Ludger Woessmann and Eric A. Hanushek, 89–200. Amsterdam:North Holland.
- Hanushek, Eric A., Susanne Link, and Ludger Woessmann.** 2012. "Does school autonomy make sense everywhere? Panel estimates from PISA." *Journal of Development Economics*, 104: 212–232.
- Heckman, James J., Hidehiko Ichimura, and Petra Todd.** 1998. "Matching As An Econometric Evaluation Estimator." *Review of Economic Studies*, 65(2): 261–294.
- Hicks, John R.** 1942. "Consumers' surplus and index-numbers." *Review of Economic Studies*, 9(2): 126–137.
- Hilger, James, Greg Rafert, and Sofia Villas-Boas.** 2011. "Expert opinion and the demand for experience goods: an experimental approach in the retail wine market." *Review of Economics and Statistics*, 93(4): 1289–1296.
- Hoffman, Donna L., and Thomas P. Novak.** 1998. "Bridging the Racial Divide on the Internet." *Science*, 280(5362): 390–391.
- Holzer, Harry J.** 1988. "Search Method Use by Unemployed Youth." *Journal of Labor Economics*, 6(1): 1–20.
- Hornung, Erik.** 2012. "Railroads and Micro-regional Growth in Prussia." *ifo Working Papers*, 127.
- IAB Europe, and McKinsey.** 2010. "Consumers driving the digital uptake."

- Imbens, Guido W., and Joshua D. Angrist.** 1994. "Identification and estimation of local average treatment effects." *Econometrica*, 62(2): 467–475.
- Jarvis, Andy, Hannes I. Reuter, Andy Nelson, and E. Guevara.** 2008. "Hole-filled SRTM for the globe, Version 4. CGIAR-CSI SRTM 90m Database." *International Center for Tropical Agriculture, Cali, Columbia*.
- Jenks, Leland H.** 1944. "Railroads as an economic force in American development." *Journal of Economic History*, 4(1): 1–20.
- Jin, Ginger Z., and Andrew Kato.** 2006. "Price, quality, and reputation: evidence from an online field experiment." *RAND Journal of Economics*, 37(4): 983–1005.
- Kang, Jun S., Polina Kuznetsova, Michael Luca, and Yejin Choi.** 2013. "Where Not to Eat? Improving Public Policy by Predicting Hygiene Inspections Using Online Reviews." *mimeo*.
- Klemens, Guy.** 2010. *The Cellphone: The History and Technology of the Gadget That Changed the World*. Jefferson, N.C.:McFarland & Company.
- Kolko, Jed.** 2012. "Broadband and local growth." *Journal of Urban Economics*, 71(1): 100–113.
- Konrad, Franz-Michael.** 2007. *Geschichte der Schule: Von der Antike bis zur Gegenwart*. München:C.H. Beck.
- Kroft, Kory, and Devin G. Pope.** 2014. "Does Online Search Crowd Out Traditional Search and Improve Matching Efficiency? Evidence from Craigslist." *Journal of Labor Economics*, 32(2): 259–303.
- Krueger, Alan B.** 2000a. "The Digital Divide in Educating African-American Students and Workers." *Princeton University Industrial Relations Section Working Paper*, 434.
- Krueger, Alan B.** 2000b. "The Internet is lowering the cost of advertising and searching for jobs." *New York Times*, July 20.
- Krueger, Alan B.** 2003. "Economic Considerations and Class Size." *Economic Journal*, 113(485): F34–F63.
- Kuhn, Peter J.** 2003. "The Internet and Matching in Labor Markets." In *New Economy Handbook*, ed. Derek C. Jones, 508–523. Amsterdam:Academic Press.
- Kuhn, Peter J., and Hani Mansour.** 2011. "Is Internet Job Search Still Ineffective?" *IZA Discussion Paper*, 5955.
- Kuhn, Peter J, and M. Skuterud.** 2004. "Internet Job Search and Unemployment Durations." *American Economic Review*, 94(1): 218–232.
- Lavy, Victor.** 2010. "Do Differences in Schools' Instruction Time Explain International Achievement Gaps? Evidence from Developed and Developing Countries." *NBER Working Paper*, 16227.
- Leuven, Edwin, and Mikael Lindahl.** 2007. "The effect of extra funding for disadvantaged pupils on achievement." *Review of Economics and Statistics*, 89(4): 721–736.
- Linden, Leigh.** 2008. "Complement or Substitute? The Effect of Technology on Student Achievement in India." *mimeo*.

- Loury, Linda D.** 2006. "Some Contacts Are More Equal than Others: Informal Networks, Job Tenure, and Wages." *Journal of Labor Economics*, 24(2): 299–318.
- Luca, Michael.** 2011. "Reviews, reputation, and revenue: The case of Yelp. com." *Harvard Business School Working Papers*, 12-016.
- Luca, Michael.** 2013. "Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud." *mimeo*.
- Luca, Michael, and Jonathan Smith.** 2011. "Salience in Quality Disclosure: Evidence from the U.S. News College Rankings." *Harvard Business School Working Papers*, 12(014).
- Lucking-Reiley, David, Doug Bryan, Naghi Prasad, and Daniel Reeves.** 2007. "Pennies from Ebay: The Determinants of Price in Online Auctions." *Journal of Industrial Economics*, 55(2): 223–234.
- Machin, Stephen, Sandra McNally, and Olmo Silva.** 2007. "New Technology in Schools: Is There a Payoff?" *Economic Journal*, 117(522): 1145–1167.
- Malamud, Ofer, and Cristian Pop-Eleches.** 2011. "Home Computer Use and the Development of Human Capital." *Quarterly Journal of Economics*, 126(2): 987–1027.
- Mang, Constantin.** 2012. "Online Job Search and Matching Quality." *ifo Working Papers*, 147.
- Mang, Constantin.** 2014. "From Upgrade to Uptake? How Mobile Internet Infrastructure Affects Local Online Services." *mimeo*.
- Mayzlin, Dina, Yaniv Dover, and Judith A. Chevalier.** 2012. "Promotional Reviews: An Empirical Investigation of Online Review Manipulation." *NBER Working Paper*, 18340.
- Metzler, Johannes, and Ludger Woessmann.** 2012. "The impact of teacher subject knowledge on student achievement: Evidence from within-teacher within-student variation." *Journal of Development Economics*, 99(2): 486–496.
- Mincer, Jacob, and Solomon Polachek.** 1974. "Family Investments in Human Capital: Earnings of Women." *Journal of Political Economy*, 82(2): 76–108.
- Moretti, Enrico.** 2011. "Social Learning and Peer Effects in Consumption: Evidence from Movie Sales." *Review of Economic Studies*, 78(1): 356–393.
- Morton, Fiona S.** 2006. "Consumer Benefit from Use of the Internet." In *Innovation Policy and the Economy, Volume 6*. Vol. 6, , ed. Adam B. Jaffe, Josh Lerner and Scott Stern, 67–90. The MIT Press.
- Morton, Fiona S., Florian Zettelmeyer, and Jorge Silva-Risso.** 2001. "Internet car retailing." *Journal of Industrial Economics*, 49(4): 501–519.
- Munnell, Alicia H.** 1992. "Infrastructure Investment and Economic Growth." *Journal of Economic Perspectives*, 6(4): 189–198.
- Nesterko, Sergiy O., Daniel T. Seaton, Konstantin Kashin, Qiuyi Han, Justin Reich, Jim Waldo, Isaac Chuang, and Andrew Dean Ho.** 2014. *HarvardX Insights Enrollment Data*.
- Nolen, Patrick J., and Stefan Klonner.** 2010. "Cell Phones and Rural Labor Markets: Evidence from South Africa." *Proceedings of the German Development Economics Conference, Hannover 2010, No. 56*.

- OECD.** 2013. *OECD Communications Outlook 2013. OECD Communications Outlook*, OECD Publishing.
- Oppenheimer, Todd.** 2003. *The flickering mind: The false promise of technology in the classroom, and how learning can be saved*. New York:Random House.
- Parmeter, Christopher F., and Jaren C. Pope.** 2012. "Quasi-Experiments and Hedonic Property Value Methods." In *Handbook on Experimental Economics and the Environment.* , ed. Michael K. Price and John A. List, 3–66. Northampton, MA:Edward Elgar Publishing.
- Peterson, Paul E.** 2010. *Saving Schools: From Horace Mann to Virtual Learning*. Cambridge, MA:Belknap Press of Harvard University Press.
- Pollakowski, Henry O.** 1995. "Data sources for measuring house price changes." *Journal of Housing Research*, 6(3): 377–387.
- Pope, Jaren C.** 2008a. "Buyer information and the hedonic: The impact of a seller disclosure on the implicit price for airport noise." *Journal of Urban Economics*, 63(2): 498–516.
- Pope, Jaren C.** 2008b. "Do Seller Disclosures Affect Property Values? Buyer Information and the Hedonic Model." *Land Economics*, 84(4): 551–572.
- Pope, Jaren C.** 2008c. "Fear of crime and housing prices: Household reactions to sex offender registries." *Journal of Urban Economics*, 64(3): 601–614.
- Resnick, Paul, Richard Zeckhauser, John Swanson, and Kate Lockwood.** 2006. "The value of reputation on eBay: A controlled experiment." *Experimental Economics*, 9(2): 79–101.
- Ridker, Ronald G., and John A. Henning.** 1967. "The determinants of residential property values with special reference to air pollution." *Review of Economics and Statistics*, 49(2): 246–257.
- Rivkin, Steven G., Eric A. Hanushek, and John F. Kain.** 2005. "Teachers, schools, and academic achievement." *Econometrica*, 73(2): 417–458.
- Rosenbaum, Paul R., and Donald B. Rubin.** 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrika*, 70(1): 41–55.
- Rosen, Sherwin.** 1974. "Hedonic prices and implicit markets: product differentiation in pure competition." *Journal of Political Economy*, 82(1): 34–55.
- Rosston, Gregory, Scott J. Savage, and Donald M. Waldman.** 2010. "Household Demand for Broadband Internet Service." *Final report to the Broadband.gov Task Force Federal Communications Commission*.
- Rouse, Cecilia E., and Alan B. Krueger.** 2004. "Putting computerized instruction to the test: a randomized evaluation of a scientifically based reading program." *Economics of Education Review*, 23(4): 323–338.
- Savage, Scott J., and Donald M. Waldman.** 2004. "United States Demand for Internet Access." *Review of Network Economics*, 3(3): 228–247.
- Schlotter, Martin, Guido Schwerdt, and Ludger Woessmann.** 2011. "Econometric methods for causal evaluation of education policies and practices: a non-technical guide." *Education Economics*, 19(2): 109–137.

- Schmitt, John, and Jonathan Wadsworth.** 2006. "Is there an impact of household computer ownership on children's educational attainment in Britain?" *Economics of Education Review*, 25(6): 659–673.
- Schultz, Christian.** 2004. "Market transparency and product differentiation." *Economics Letters*, 83(2): 173–178.
- Schwerdt, Guido, and Amelie C. Wuppermann.** 2011. "Economics of Education Review Is traditional teaching really all that bad? A within-student between-subject approach." *Economics of Education Review*, 30(2): 365–379.
- Simon, Curtis J., and John T. Warner.** 1992. "Matchmaker, Matchmaker: The Effect of Old Boy Networks on Job Match Quality, Earnings, and Tenure." *Journal of Labor Economics*, 10(3): 306–330.
- Solow, Robert M.** 1987. "We'd Better Watch Out." *New York Times*, July 12.
- Stevenson, Betsey.** 2006. "The Impact of the Internet on Worker Flows." *Unpublished paper*.
- Strömberg, David.** 2004. "Radio's impact on public spending." *Quarterly Journal of Economics*, 119(1): 189–221.
- Terracciano, Antonio, Robert R. McCrae, and Paul T. Costa.** 2010. "Intra-individual Change in Personality Stability and Age." *Journal of Research in Personality*, 44(1): 31–37.
- The Economist.** 2012. "Local services online: Your friendly neighbourhood app." *The Economist*, October 27th 2012.
- Theunissen, Gert, Marijke Verbruggen, Anneleen Forrier, and Luc Sels.** 2011. "Career sidestep, wage stepback? The impact of different types of career breaks on wages." *Gender, Work & Organization*, 18(1): 110–131.
- TÜVRheinland.** 2013. "Szenarien und Kosten für eine kosteneffiziente flächendeckende Versorgung der bislang noch nicht mit mindestens 50 Mbit / s versorgten Regionen."
- Umweltausschuss des Bundestages.** 2001. "Oeffentliche Anhoerung Mobilfunk am 2. Juli 2001." *Ausschussdrucksache 14 / 577*.
- van Ours, Jan C., and Milan Vodopivec.** 2008. "Does reducing unemployment insurance generosity reduce job match quality?" *Journal of Public Economics*, 92(3-4): 684–695.
- Varian, Hal R.** 1980. "A model of sales." *American Economic Review*, 70(4): 651–659.
- Varian, Hal R.** 1982. "The Nonparametric Approach to Demand Analysis." *Econometrica*, 50(4): 945–973.
- Varian, Hal R.** 2009. "Online Ad Auctions." *American Economic Review*, 99(2): 430–434.
- Varian, Hal R.** 2010. "Computer Mediated Transactions." *American Economic Review: Papers & Proceedings*, 100(May): 1–10.
- Varian, Hal R.** 2011. "Economic Value of Google." *Presentation at Web 2.0 Summit*, March 29.
- Varian, Hal R.** 2012. "Revealed Preference and its Applications." *Economic Journal*, 122(560): 332–338.
- Vermeulen, Ivar E., and Daphne Seegers.** 2009. "Tried and tested: The impact of online hotel reviews on consumer consideration." *Tourism Management*, 30(1): 123–127.



- Waldrop, Mitchell.** 2013. "Campus 2.0." *Nature*, 495: 160–163.
- Waverman, Leonard, and Lars-Hendrik Röller.** 2001. "Telecommunications infrastructure and economic development: a simultaneous approach." *American Economic Review*, 91(4): 909–923.
- Weitzman, ML.** 1998. "Recombinant growth." *Quarterly Journal of Economics*, 113(2): 331–360.
- West, Joel, and Michael Mace.** 2010. "Browsing as the killer app: Explaining the rapid success of Apple's iPhone." *Telecommunications Policy*, 34(5-6): 270–286.
- Williams, Donald R.** 2000. "Consequences of Self-Employment for Women and Men in the United States." *Labour Economics*, 7(5): 665–687.
- Woessmann, Ludger.** 2003. "Schooling resources, educational institutions and student performance: the international evidence." *Oxford Bulletin of Economics and Statistics*, 65(2): 117–170.
- Ye, Qiang, Rob Law, Bin Gu, and Wei Chen.** 2011. "The influence of user-generated content on traveler behavior: An empirical investigation on the effects of e-word-of-mouth to hotel online bookings." *Computers in Human Behavior*, 27(2): 634–639.
- Zweifel, Peter, and Christoph Zaborowski.** 1996. "Employment Service: Public or Private?" *Public Choice*, 89(1): 131–162.



# CURRICULUM VITAE

CONSTANTIN MANG

BORN 06/11/1986 IN MUNICH, GERMANY

2010 – 2014 IFO INSTITUTE - LEIBNIZ-INSTITUTE FOR ECONOMIC RESEARCH  
AT THE UNIVERSITY OF MUNICH

*Junior Economist and PhD Candidate*

2009 – 2010 WHU – OTTO BEISHEIM SCHOOL OF MANAGEMENT, VALLENDAR  
AND BUCERIUS LAW SCHOOL, HAMBURG

*Master of Law and Business (2010)*

2006 – 2009 UNIVERSITÀ COMMERCIALE LUIGI BOCCONI, MILAN

*Bachelor of International Economics and Management (2009)*

2004 – 2006 SALEM INTERNATIONAL COLLEGE, ÜBERLINGEN

*Abitur and IB Certificate in Economics (2006)*

1997 – 2004 KREUZBURG GYMNASIUM, GROSSKROTZENBURG