



Essays in Applied Microeconomics

Ingvil Gaarder

Thesis submitted for assessment with a view to obtaining the degree
of Doctor of Economics of the European University Institute

Florence, September 2014

European University Institute
Department of Economics

Essays in Applied Microeconomics

Ingvil Gaarder

Thesis submitted for assessment with a view to obtaining the degree of Doctor of Economics of the European University Institute

Examining Board

Prof. Jérôme Adda, EUI & Bocconi University, Supervisor

Prof. Andrea Ichino, EUI

Prof. Kjell G. Salvanes, Norwegian School of Economics

Prof. Guy Michaels, LSE

© Gaarder, 2014

No part of this thesis may be copied, reproduced or transmitted without prior permission of the author

Abstract

The first chapter of this thesis examines the incidence and distribution effects of VAT. The context of the study is a sharp change in the VAT policy on food items in Norway. I examine the direct impact of the policy change on the consumer prices of food items as well as any cross-price effects on other goods. The estimates suggest that taxes levied on food items are completely shifted to consumer prices, whereas the pricing of other goods is not materially affected. To understand the distributional effects of the VAT reform, I use expenditure data and estimate the compensating variation of the tax induced price changes. I find that lowering the VAT on food attenuates inequality in consumer welfare.

Chapter 2 and 3 investigates how technology adoption in firms affects productivity and labor market outcomes. In chapter 2 (joint work with A. Akerman and M. Mogstad), the focus is on whether adoption of broadband internet in firms enhance labor productivity and increase wages. We find that broadband adoption favors skilled labor by increasing its relative productivity. By comparison, broadband internet is a substitute for workers without high school diploma, lowering their marginal productivity. Consistent with the estimated changes in labor productivity, wage regressions show the expansion of broadband internet improves (worsens) the labor outcomes of skilled (unskilled) workers. We explore several possible explanations for the skill bias of broadband internet. We find suggestive evidence that broadband internet complements skilled workers in executing nonroutine abstract tasks, and substitutes for unskilled workers in performing routine tasks.

In chapter 3, I move attention to why firms differ so much in their abilities to convert inputs into output. One possibly explanation is that firms choose different technology, giving variation in output that cannot be explained based on observable inputs like standard labor or capital measures. I find that broadband internet accounts for about 2 to 3 percent of the standard deviation in TFP across firms. I decompose aggregate TFP in the economy into an unweighted average of firm-level TFP and the covariance between market share and firm-level TFP. I find that the contribution of the covariance term to aggregate TFP is reduced substantially when I account for the variation in output that can be explained by broadband adoption in firms. This suggests that broadband internet accounts for some of the allocation of production to more productive firms.

Acknowledgments

First and foremost, I would like to thank my advisor, Jérôme Adda. This thesis has benefited greatly from his comments and suggestions. I would also like to thank Anders Akerman and Magne Mogstad. I've learned a lot from working with the two of you. In addition, I would like to express my gratitude towards Magne for his constant encouragement and support during these years.

Parts of this thesis has been written during a research stay at the Centre for Microdata Methods and Practice (Cemmap). I would especially like to thank Lars Nesheim for making the stay possible. I benefited a great deal from interacting with the research group both at Cemmap and the IFS.

I also wish to thank the Norwegian Research Council, Norges Bank and the EUI for financial support.

Contents

Preface	1
1 Incidence and Distributional Effects of Value Added Taxes	3
1.1 Introduction	3
1.2 Data and Background	5
1.2.1 Data sources and summary statistics	5
1.2.2 The VAT reform and institutional details	8
1.2.3 Expected reform effects	9
1.3 Incidence of the VAT reform	11
1.3.1 Research design	11
1.3.2 Graphical evidence	11
1.3.3 Regression estimates	14
1.3.4 Specification checks	15
1.4 Distributional effects of VAT reform	17
1.4.1 Model and estimates of demand system	17
1.4.2 Distributional effects	20
1.5 Conclusion	22
2 The Skill Complementarity of Broadband Internet	25
2.1 Introduction	25
2.2 Technological change in production	28
2.2.1 A framework for estimation	28
2.2.2 Empirical challenges and previous research	29
2.3 Data	30
2.4 Expansion of broadband internet	33
2.5 Empirical strategy	34
2.6 Empirical results	37
2.6.1 Evidence on skill bias	37
2.6.2 Specification checks	39
2.6.3 Understanding the pattern of adoption	46
2.6.4 Economic significance	48
2.6.5 Worker heterogeneity	48
2.7 Mechanisms	49
2.8 Conclusion	55

2.A Appendix A: Additional Tables and Figures	56
2.B Appendix B: Estimation Details	70
3 Technology Adoption and Productivity Dispersion: Evidence from Broadband Adoption in Firms	73
3.1 Introduction	73
3.2 Productivity differences with technology choice	76
3.2.1 A framework for estimation	76
3.2.2 Empirical challenges and previous research	76
3.3 Data	77
3.4 Expansion of broadband internet	80
3.5 Technological change in production	81
3.5.1 Identification of production functions	81
3.5.2 Estimated production functions	85
3.6 Implications of broadband internet for TFP	90
3.6.1 Dispersion of TFP across firms	90
3.6.2 Aggregate TFP	90
3.7 Conclusion	91
3.A Appendix A: Additional Tables and Figures	93
3.B Appendix B: Estimation Details	99

Preface

Much of the controversy surrounding recent policy proposals to broaden the base for value added taxes (VAT) revolves around who ultimately bears the burden of these taxes. The typical assumption is that consumer prices fully reflect taxes, so that the only important empirical question is how the tax induced price changes affect members of different income groups. However, the evidence base is scarce and market imperfections could generate both over and undershifting of VAT to consumer prices. In the first chapter, I examine the incidence and distributional effects of VAT in a setting with plausibly exogenous variation in tax rates. The context of the study is a sharp change in the VAT policy on food items in Norway. Using a regression discontinuity design, I examine the direct impact of the policy change on the consumer prices of food items as well as any cross-price effects on other goods. The estimates suggest that taxes levied on food items are completely shifted to consumer prices, whereas the pricing of other goods is not materially affected. To understand the distributional effects of the VAT reform, I use expenditure data and estimate the compensating variation of the tax induced price changes. I find that lowering the VAT on food attenuates inequality in consumer welfare, in part because households adjust their spending patterns in response to the prices changes. By comparison, the usual first order approximation of the distributional effects, which ignores behavioural responses, seriously understates the redistributive nature of the VAT reform.

In the second chapter (joint work with A. Akerman and M. Mogstad), we exploit rich Norwegian data with firm-level information on value added, factor inputs and broadband adoption to answer the questions: Does adoption of broadband internet in firms enhance labor productivity and increase wages? And is this technological change skill biased or factor neutral? We estimate production functions where firms can change their technology by adopting broadband internet. A public program with limited funding rolled out broadband access points, and provides plausibly exogenous variation of broadband adoption in firms. This enables us to address endogeneity of broadband adoption and examine how it shifts the production technology and changes the productivity and labor outcomes of different types of workers. We find that broadband adoption favors skilled labor by increasing its relative productivity. The increase in productivity of skilled labor is especially large for college graduates in fields such as science, technology, engineering and business. By comparison, broadband internet is a substitute for workers without high school diploma, lowering their marginal productivity. Consistent with the estimated changes in labor productivity, wage regressions show the expansion of broadband internet improves (worsens) the labor outcomes of skilled (unskilled) workers. We explore several possible explanations for the skill bias of broadband internet. We find suggestive evidence that broadband internet complements skilled workers in executing nonroutine abstract tasks, and substitutes for unskilled workers in performing routine tasks. Taken together, our findings have important implications for the ongoing policy debate over government investment in broadband infrastructure to encourage productivity and wage growth.

In the last chapter, I explore why firms differ so much in their abilities to convert inputs into output.

One possible explanation is that firms choose different technology, giving variation in output that cannot be explained based on observable inputs like standard labor or capital measures. In this chapter, I estimate how the adoption of new technology induces productivity changes and explains dispersion in total factor productivity (TFP) across firms. As in chapter 2, the context is the adoption of broadband internet in Norwegian firms over the period 2001-2007. I begin by estimating a standard production function where any output driven by broadband internet will show up as TFP. Next, I estimate a production function where firms can change their technology by adopting broadband internet. I allow for the possibility that broadband adoption is factor-neutral or factor-biased, and therefore, directly or indirectly affect output. I rely on the same public program, as in chapter 2, as a source of plausible exogenous variation in broadband adoption. I find that broadband internet accounts for about 2 to 3 percent of the standard deviation in TFP across firms. I decompose aggregate TFP in the economy into an unweighted average of firm-level TFP and the covariance between market share and firm-level TFP. I find that the contribution of the covariance term to aggregate TFP is reduced substantially when we account for the variation in output that can be explained by broadband adoption in firms. This suggests that broadband internet accounts for some of the allocation of production to more productive firms.

Chapter 1

Incidence and Distributional Effects of Value Added Taxes

1.1 Introduction

Are taxes levied on commodities completely shifted to consumer prices, or does the incidence also fall on firms? What are the welfare implications of commodity taxes for poor and rich households? These questions are important for both policy and scientific research, as commodity taxes make up a large part of fiscal revenue in most developed countries. In Europe, much of the controversy surrounding recent policy proposals to broaden the base for value added taxes (VAT) revolves around who ultimately bears the burden of these taxes. In the United States, recent debates on whether to increase reliance on consumption-based taxes have raised concerns over the distributional effects of such policy changes. The typical assumption is that consumer prices fully reflect taxes, so that the only important empirical question is how the tax induced price changes affect members of different income groups. However, the evidence base is scarce (Crawford, Keen, and Smith, 2010) and market imperfections could generate both over and undershifting of commodity taxes to consumer prices (Seade, 1987; Delipalla and Keen, 1992; Anderson, De Palma, and Kreider, 2001).

The aim of this paper is to empirically investigate the incidence and distributional effects of commodity taxes in a setting with plausibly exogenous variation in tax rates. The context of our study is an abrupt change in the VAT policy on food in Norway. As in most European countries, food retailing in Norway is highly concentrated with a few chains commanding most of the market. On July 1st, 2001 the Norwegian government reduced the VAT on all food items from 24 to 12 percent, while the VAT on non-food items remained at 24 percent. This sharp change in VAT policy provides an attractive setting to analyze the passthrough of commodity taxes using a regression discontinuity (RD) design that compares consumer prices just before (i.e. the control group) and after (i.e. the treatment group) the reform date. We apply this design to rich data on retail prices for a representative sample of consumer goods. This allows us to estimate the direct impact of the policy change on the consumer prices of food items as well as any cross-price effects on other goods. We challenge the identifying assumptions of the RD design through a number of robustness checks, finding little cause for worry.

The RD estimates tell us whether the gains from the VAT reform ultimately fall on consumers or producers. However, the distributional effects also depend on the extent to which poor and rich households are affected by the passthrough to consumer prices. Using survey data on consumer expenditure, we perform a

first order approximation of the distributional effects. An advantage of this approach is that it simply requires information on the price changes and the pre-reform expenditure patterns of the households. However, the VAT reform generated substantial rather than marginal changes in food prices. In such cases, substitution effects can be non-trivial, as consumers substitute towards relatively cheaper goods. The first order approximations ignore these effects, and therefore, can be seriously biased (see e.g. Banks, Blundell, and Lewbel, 1996). To address this concern, we use expenditure data to estimate the Almost Ideal (AI) demand system. This allows us to incorporate behavioural responses in estimating the compensating variation of the changes in prices associated with the VAT reform.

The insights from our empirical results may be summarized with two broad conclusions. First, the VAT on food items are completely shifted to consumer prices, implying that producers bear little, if anything, of the tax burden. By comparison, there is little if any spillover effects of the VAT reform to the consumer prices of non-food items. Second, lowering the VAT on food substantially attenuates inequality in consumer welfare. This reduction in inequality is partly because poor households have a higher expenditure share on food prior to the reform, but also because households adjust their spending patterns in response to price changes. By comparison, the usual first order approximation of the distributional effects, which ignores behavioural responses, seriously understates the redistributive nature of the VAT reform.

Our findings have implications for recent proposals for tax reforms. For example, the Mirrlees Review (2012) sets out a comprehensive proposal for tax reform in the United Kingdom. A key element of the reform package is to broaden the base for VAT, in part by removing the zero rating for food. Arguing that there is little hard evidence to draw on, the Mirrlees Review assumes the incidence of VAT is fully on consumer prices. Atkinson (2013) questions the reform proposal, stressing that until direct evidence is available, “we should remain agnostic about the strength of the optimal tax argument for extending VAT to food” (p. 6). Our paper helps to close that gap by providing transparent and credible identification of the incidence of VAT taxes on food.

Our paper is primarily related to two recent empirical studies on the incidence of VAT. Carbonnier (2007) studies two VAT reforms in France which reduced the rates on new car sales and on housing repair services. He uses the variation in consumer prices across goods and over time to estimate the passthrough of these VAT reforms. His estimates suggest a majority of the tax burden is paid by consumers, especially in the competitive market for housing repair services. Kosonen (2013) analyze the incidence of VAT in the context of hairdressing services in Finland. He uses a difference-in-differences strategy where the control group consists of beauty salons, day spas and massage services. His estimates suggest the tax burden on hairdressing services is shared between consumers and producers.

Our paper expands on this research in several important ways. First, our study provides novel evidence on the incidence of a VAT tax system with lower rates on perceived necessities such as food. Second, our rich data allows us to look at cross-price effects on other goods, and therefore, capture the entire change in the price structure. Third, we quantify the extent to which a lower VAT rate on food redistributes resources from better-off households to less well-off households. To the best of our knowledge, substitution effects have not been incorporated in distributional analysis of VAT reforms. We do so here, and investigate the accuracy of the usual first order approximation to welfare implications of tax reform.

Our paper is also related to a sparse empirical literature on the passthrough of sales taxes. Unlike VAT, a sales tax is imposed only at the retail level, which could have important implications for how the tax burden is shared between consumers and producers (see e.g. Anderson, De Palma, and Kreider, 2001). Poterba (1996) and Besley and Rosen (1999) examined tax shifting in the United States by comparing local sales

taxes and consumer prices across areas and over time. Their estimates suggest that consumers tend to pay for sales taxes. In addition, researchers have examined the incidence of per unit (excise) taxes on goods such as tobacco, alcoholic beverages, and gasoline.¹ Economic theory predicts that in markets with imperfect competition, the consumer share of excise taxes could differ from that of VAT. This theoretical prediction is supported by the empirical evidence in Delipalla and O'Donnell (2001) and Carbonnier (2013).

The remainder of the paper proceeds as follows. In Section 2, we describe our data and discuss the VAT reform and its expected impact. In Section 3, we discuss the RD design, present our main findings on VAT incidence, and reports robustness checks. In Section 4, we present the demand system and analyze the distributional effects of VAT. Section 5 concludes.

1.2 Data and Background

1.2.1 Data sources and summary statistics

Our analysis uses two data sources. The first is a rich data set on retail prices for a representative sample of consumer goods. The data is collected by Statistics Norway and forms the basis for calculating the Norwegian Consumer Price Index. Every month, Statistics Norway collects information about the consumer prices on a variety of items. The reported consumer price is the current price on the 15th of each month and the retailers report this price the first working day after the 15th. In 2001, there was more than 180,000 recorded prices from a total of 545 different retailers. Because of the detailed nature of the data, it is possible to follow prices on a given item and retailer over time.

Table 1.1 displays summary statistics for the major consumer goods in 2001. This table shows that the food category consist of 250 different items. In each category, the average consumer price is computed as the weighted average of the retail prices on the items that belong to this category. We follow the procedure used to construct the Consumer Price Index in the choice of weights and classification of items. For food, the average consumer price is about 30 NOK per item. In general, we see that the consumer prices vary considerably between the different types of goods.

The other data source we will be using is the Norwegian Consumer Expenditure Survey for the years 1989–2001. In addition to detailed information on each household expenditure, there is also a rich set of household characteristics, including information on household size, age of the household members, gender, marital status, region, labor status, occupation, and household disposable income. We use the same classification of goods for the expenditure data as for the price data. Following Blundell, Pashardes, and Weber (1993), our sample consists of households in which the household head is between 20 and 60 years old and not self-employed; the sample is top and bottom coded at the the 99th and 1st percentile level of the distribution of household income.² Throughout the paper, we use sampling weights to produce representative estimates for the corresponding population of households.

Table 1.2 summarizes the expenditure shares for non-durable goods. As expected, food purchases form the largest share of household expenditure, and the expenditure share declines in household income. For example, food purchases make up 28.3 % of household expenditure in the bottom quartile of the household

¹For example, Doyle and Samphantharak (2008) and Marion and Muehlegger (2011) study tax incidence on gasoline. Their findings point to a complete passthrough to consumer prices. Other studies have examined the consumer share of excise taxes on tobacco and alcohol beverages (see e.g. Young and Bielińska-Kwapisz, 2002, Kenkel, 2005 and DeCicca, Kenkel, and Liu, 2013).

²This selection is meant to remove the possibility that small outliers in the income distribution is creating nonlinearities in the budget-share equations.

Table 1.1: Summary statistics: Consumer price data

	Average	No. of		
	Consumer Price	Items	Retailers	Obs.
Food	29.8	250	545	180,510
Clothing	391	100	445	41,303
Services	295	54	615	29,419
Fuel	37.9	6	328	1,988
Alcohol	13.8	3	277	4,430
Transport	898	28	377	33,493
Other non-durables	207	217	1,198	118,284

Notes: This table displays summary statistics for non-durable goods in 2001. In each category, the average consumer price is computed as the weighted average of the retail prices on the items that belong to this category. We follow the procedure used to construct the Consumer Price Index in the choice of weights and classification of items.


income distributions, whereas the expenditure share on food is only 23.9 % in the top quartile. We see the same pattern for other perceived necessities such as fuel, while the share of household expenditure on goods like clothing and transport increases in household income.

Table 1.2: Summary statistics: Expenditure data

Expenditure shares:	Household income											
	Full Sample		1st quantile		2nd and 3rd quantile		4th quantile					
	Mean	St.dev	Mean	St.dev	Mean	St.dev	Mean	St.dev				
Food	0.268	0.12	0.283	0.14	0.276	0.12	0.239	0.11				
Clothing	0.103	0.10	0.091	0.10	0.103	0.09	0.117	0.09				
Services	0.160	0.12	0.159	0.14	0.154	0.12	0.173	0.12				
Fuel	0.087	0.06	0.111	0.08	0.084	0.05	0.070	0.04				
Alcohol	0.011	0.02	0.012	0.03	0.011	0.02	0.011	0.02				
Transport	0.144	0.13	0.129	0.14	0.146	0.12	0.154	0.13				
Other non-durables	0.227	0.14	0.214	0.15	0.227	0.14	0.238	0.14				
No. of obs	11107		2777		5554		2776					

Notes: The sample consists of households in which the household head is between 20 and 60 years old and not self-employed. This sample is top and bottom coded at the the 99th and 1st percentile level of the distribution of household income. Column 1-2 shows means and st. deviations in expenditure shares on various goods in the full sample, while columns 3-8 report the same statistics for different income groups.

Figure 1.1: Nonparametric Engel Curve for Food Shares



Notes: The solid line is estimated using a nonparametric kernel regression with Gaussian kernel and a mean integrated squared-error optimal smoothing parameter. The shaded area shows the 95 % confidence bands.

Figure 1.1 looks closer at the relationship between household income and food expenditure by graphing the Engel curve. This figure provides a nonparametric description of the Engel curve and suggests that a log-linear specification approximates well the food share curve. This result aligns well with previous evidence from developed countries (see e.g. Banks, Blundell, and Lewbel, 1997)

1.2.2 The VAT reform and institutional details

In Norway, VAT are levied on the sale of goods and services on registered businesses with annual turnover above NOK 50,000. VAT applies to all sales, whether to private consumers or other businesses. Under the “invoice-credit” form of the VAT, registered businesses offset the VAT they have been charged on purchases against the liability on their sales, remitting only the net amount due. The result is that no net revenue is collected from the taxation of intermediate goods sales, so that the ultimate base of the tax is final consumption.³

Before the VAT reform in 2001, Norway had a VAT rate of 24 % on most goods. Notable exceptions were certain transport services which had reduced rates and newspapers and books with zero ratings. On July 1st, 2001 the Norwegian government reduced the VAT on all food items from 24 to 12 percent, while the VAT on other goods did not change. The reduction in the VAT on food items was announced in December 2000.

The key motivation for the reform was that the broad-based VAT system with a uniform rate on most goods places a too large tax burden on poor households. The view that reduced VAT on food items would reduce the tax burden of the poor was based on the standard assumption that the tax would be shifted forward to consumers through price decreases. However, the market structure may affect the incidence of the tax, as discussed in greater detail below. Like most European countries, food retailing in Norway is highly concentrated. As shown in Panel A of Table 1.3, the largest chain had in 2004 a market share of 34.6 %,

³We refer to Crawford, Keen, and Smith (2010) for a detailed discussion of the structure of VAT in OECD countries.

Table 1.3: Concentration in the food retail market in Norway, 2004

<i>Panel A:</i>	Largest chain	Second largest chain	Third largest chain
Market share:	34.7 %	23.7 %	23.6 %
<i>Panel B:</i>	All food retailers		
HHI index:	2600		

Notes: Panel A shows the the market shares of the three largest chains in the food retail market. Panel B shows the Herfindahl-Herschman Index (HHI), calculated as $HHI = \sum_{i=1}^I s_i^2$, where s_i represents market share of retailer i and I is the total number of retailers. Source: Einarsson (2007).

whereas the three largest chains commanded 82 % of the market. The Herfindahl-Herschman Index (HHI) is reported in Panel B. The US regulatory authorities, like other competition authorities, use HHI in their assessment of mergers. If there is only one firm in the market, the HHI will equal 10,000; if the market is divided equally between a large number of firms the HHI will approach 0; and if the figure is higher than 1800, US law states there is a risk of significant concentration and any potential merger under such circumstances is subjected to careful scrutiny. By this standard, the retail market in Norway was highly concentrated with a HHI of 2600. The Norwegian food retail market is also highly concentrated in comparison with other countries: Einarsson (2007) reports HHI figures of 1600 in France and Germany, 1800 in the United Kingdom, and as low as 300–500 in Spain.

1.2.3 Expected reform effects

As in the Mirrles Review, it is customary to think of the burden of VAT as being borne by consumers in the form of higher after-tax prices, but in theory there is considerable scope for shifting of the tax burden. Indeed, there are plausible circumstances in which consumers bear more than 100 % of the burden or pay little if anything of the VAT. Below, we describe the measures of tax shifting we will use and briefly discuss the challenges to making informative theoretical predictions of the passthrough of the VAT reform.

Tax shifting measures. Let τ denote the VAT rate. The producer (or pre-tax) price is $\frac{p}{1+\tau}$ where p is the consumer (or after-tax) price. The amount of taxes paid per unit sold is $\frac{\tau p}{1+\tau}$. After a change in the VAT rate, the consumer price variation is $\frac{dp}{d\tau}$ and the tax variation is $\frac{d}{d\tau} \left(\frac{\tau p}{1+\tau} \right)$. The consumer share of the change in VAT is then given by

$$cs = \frac{(1+\tau)}{p} \frac{dp}{d\tau} \frac{(1+\tau)}{1+\tau \frac{(1+\tau)}{p} \frac{dp}{d\tau}}. \quad (1.1)$$

A consumer share of more than one means the tax is overshifted. If the consumer share is equal to one the tax is fully forward shifted, and consumers bear the full cost of the VAT change. By comparison a consumer share less than one implies the tax is undershifted, and producers bear some of the cost.

Equation (1.1) tells us the tax is fully forward shifted when $\frac{dp}{d\tau} = \frac{p}{1+\tau}$, which is equivalent to $\frac{d}{d\tau} \left(\frac{p}{1+\tau} \right) = 0$. This means the tax is fully shifted whenever the producer price does not change as a result of the VAT change. In our setting, this implies the VAT reduction from 24 % to 12 % on food items is fully forward shifted if consumer prices on food decreases with 10.2 %.

Theory predictions. A priori, it is challenging to credibly predict the passthrough of a VAT reform, as

it requires detailed knowledge of the market structure and reliable estimates of demand and supply.

Consider first the benchmark of perfect competition, in which case the consumer price variation is given by

$$\frac{dp}{d\tau} = \frac{p}{(1+\tau)} \frac{\eta_S}{\eta_S + \eta_D}, \quad (1.2)$$

where $\eta_S = \frac{p}{(1+\tau)} \frac{1}{S} \frac{\partial S}{\partial p}$ is the supply elasticity evaluated at the producer price and $\eta_D = -\frac{p}{D} \frac{\partial D}{\partial p}$ is the demand elasticity evaluated at the consumer price. Equation 1.2 shows that even with perfect competition, it is difficult to predict the passthrough of a VAT reform: While overshifting is not possible, the consumer share can range from 0 to a 100 %. If demand for the taxed good is relatively elastic compared to supply then producers bear most of the tax burden, whereas the consumer share is larger if demand is less elastic than supply.

In our setting, there is strong market concentration and imperfect competition is likely, which makes it even more difficult to credibly predict the passthrough of a VAT reform.⁴ To see this, suppose there are n firms and each firm produces a variant of a differentiated product. Firm i 's profit is given by

$$\pi_i = \left[\frac{p_i}{(1+\tau_i)} - c_i \right] D_i(p_i; \mathbf{p}_{-i}).$$

where c_i is a firm specific constant marginal cost, $D_i(p_i; \mathbf{p}_{-i})$ is the demand for firm i 's product as a function of firm i 's own consumer price, p_i , and a vector consisting of the other firms' consumer prices, \mathbf{p}_{-i} . Further, the function $D_i(p_i; \mathbf{p}_{-i})$ is continuously differentiable, decreasing in p_i and increasing in all elements of \mathbf{p}_{-i} . At a Bertrand-Nash equilibrium, assuming an interior solution, each firm will set a price p_i , given \mathbf{p}_{-i} , such that the first order condition is satisfied:

$$(p_i - \tilde{c}_i) \frac{\partial D_i(p_i; \mathbf{p}_{-i})}{\partial p_i} + D_i(p_i; \mathbf{p}_{-i}) = 0, \quad (1.3)$$

where $\tilde{c}_i = (1 + \tau_i)c_i$ denotes effective cost.

The effects of an increase in the VAT on own and other producer prices are given by total differentiating the first order conditions given in (1.3)

$$\frac{dp_i}{d\tau_i} = \frac{p_i}{1 + \tau_i} \frac{1 + \varepsilon_{ii}}{2\varepsilon_{ii} - E_{ii}} - \sum_{j \neq i} \frac{p_j}{p_j} \frac{\varepsilon_{ii} E_{ij} + E_{ij} - \varepsilon_{ij}}{2\varepsilon_{ii} - E_{ii}} \frac{dp_j}{d\tau_i} \quad (1.4)$$

and for $j \neq i$:

$$\frac{dp_j}{d\tau_i} = - \sum_{k \neq j} \frac{p_k}{p_k} \frac{\varepsilon_{jj} E_{jk} + E_{jk} - \varepsilon_{jk}}{2\varepsilon_{jj} - E_{jj}} \frac{dp_k}{d\tau_i} \quad (1.5)$$

where we have substituted c_i from the first order condition (1.3), $\varepsilon_{ij} = \frac{\partial D_i}{\partial p_j} \frac{p_j}{D_i}$ is the own or cross price elasticity of demand, and $E_{ij} = \frac{\partial^2 D_i}{\partial p_i \partial p_j} \frac{p_j}{\partial D_i / \partial p_i}$ is the elasticity of the slope of the demand curve.

While under perfect competition the passthrough rate is entirely determined by the elasticity of supply and demand, Equation (1.4) shows that under imperfect competition the curvature of demand also plays a central role. Consider, for example, the case in which an increase in the VAT of good i does not lead to a

⁴See Anderson, De Palma, and Kreider (2001) for a theoretical analysis of the incidence of VAT in an oligopolistic industry with differentiated products and price-setting (Bertrand) firms. Seade (1987) and Delipalla and Keen (1992) provide a theoretical analysis of incidence in the case of a oligopolistic industry with homogenous demand and quantity-setting (Cournot) firms.

price change for good j , $\frac{dp_j}{ds_i} = 0 \forall j \neq i$. In this case, equation (1.4) is equal to

$$\frac{dp_i}{d\tau_i} = \frac{p_i}{(1 + \tau_i)} \frac{1 + \varepsilon_{ii}}{2\varepsilon_{ii} - E_{ii}}.$$

From the consumer share equation (1.1), it follows that the consumer share exceeds a 100 % if the curvature of the demand function is such that $E_{ii} > \varepsilon_{ii} - 1$. Because standard demand forms restrict this curvature in ways that have little empirical or theoretical foundation (see e.g. Fabinger and Weyl, 2012), imperfect competition makes it particularly difficult to credibly predict the passthrough rate.

Taken together, the challenges of making informative theoretical predictions motivate our empirical analysis of the incidence of VAT in a setting with plausibly exogenous variation in tax rates.

1.3 Incidence of the VAT reform

1.3.1 Research design

On July 1st, 2001 the VAT on all food items was reduced from 24 to 12 percent, while the VAT on non-food items remained at 24 percent. This sharp change in VAT policy provides an attractive setting to analyze the passthrough of commodity taxes using a RD design that compares consumer prices just before (i.e. the control group) and after (i.e. the treatment group) the reform date.

Our RD design can be described by the following regression model:⁵

$$y_{it} = \alpha + \mathbf{1}\{t \geq c\} [g_l(t - c) + \lambda] + \mathbf{1}\{t < c\} g_r(c - t) + e_{it} \quad (1.6)$$

where y_{it} denotes log consumer price on good i in month t , c is the reform date (July 1st, 2001), e_{it} is an error term, and g_l , and g_r are unknown functions. The key identifying assumptions is that prices do not change in anticipation of the VAT reform and that other factors determining consumer prices evolves smoothly around the reform date. Under these assumptions, we can consistently estimate the parameter λ , which gives the impact of the VAT reform on the consumer price of good i . Below, we challenge the identifying assumptions of the RD design, finding little cause for worry.

To implement the RD design, we need to specify g_l and g_r and decide on the window on each side of the reform date. Our first specification uses a local linear regression with triangular kernel density and 2 months of bandwidth on each side of the reform date. Our second specification uses a window of just one month on each side of the reform date. Because we have monthly data on consumer prices, the RD model is then equivalent to a first-difference (FD) model: the average consumer prices in June 2001 is compared to the average consumer prices in July 2001.⁶

1.3.2 Graphical evidence

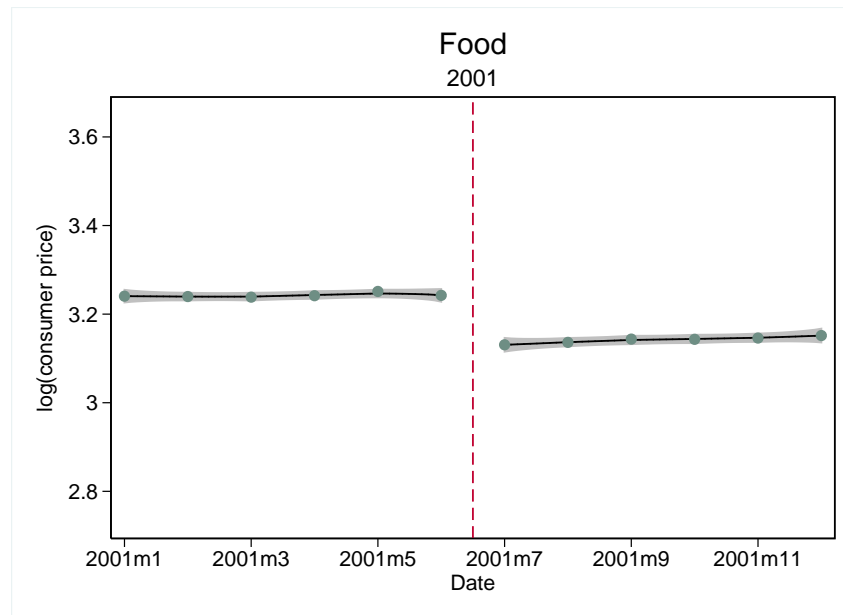
A virtue of the RD design is that it provides a transparent way of showing how the reform impact is identified. To this end, we begin with a graphical depiction before turning to a more detailed regression-based analysis.

Figure 1.2 shows both the unrestricted and the estimated monthly means of consumer prices for food items during 2001. The estimated monthly means come from a local linear regression with a triangular

⁵See e.g. Lee and Lemieux (2009) for a detailed discussion of the RD design.

⁶When following the procedure of Imbens and Kalyanaraman (2012), we estimate an optimal bandwidth which of 1.37 months. This result motivates the choice of bandwidths in the FD and RD specifications.

Figure 1.2: Evolution of consumer price on food over time



Notes: Each observation is the average consumer price for food reported on the 15th each month. The dashed vertical line denote the reform date. The solid lines are from a local linear regression with triangular weights on monthly consumer price data. The shaded area shows the 95 % confidence bands. The y-axes are scaled to $\pm .5$ st.dev. of the mean consumer price.

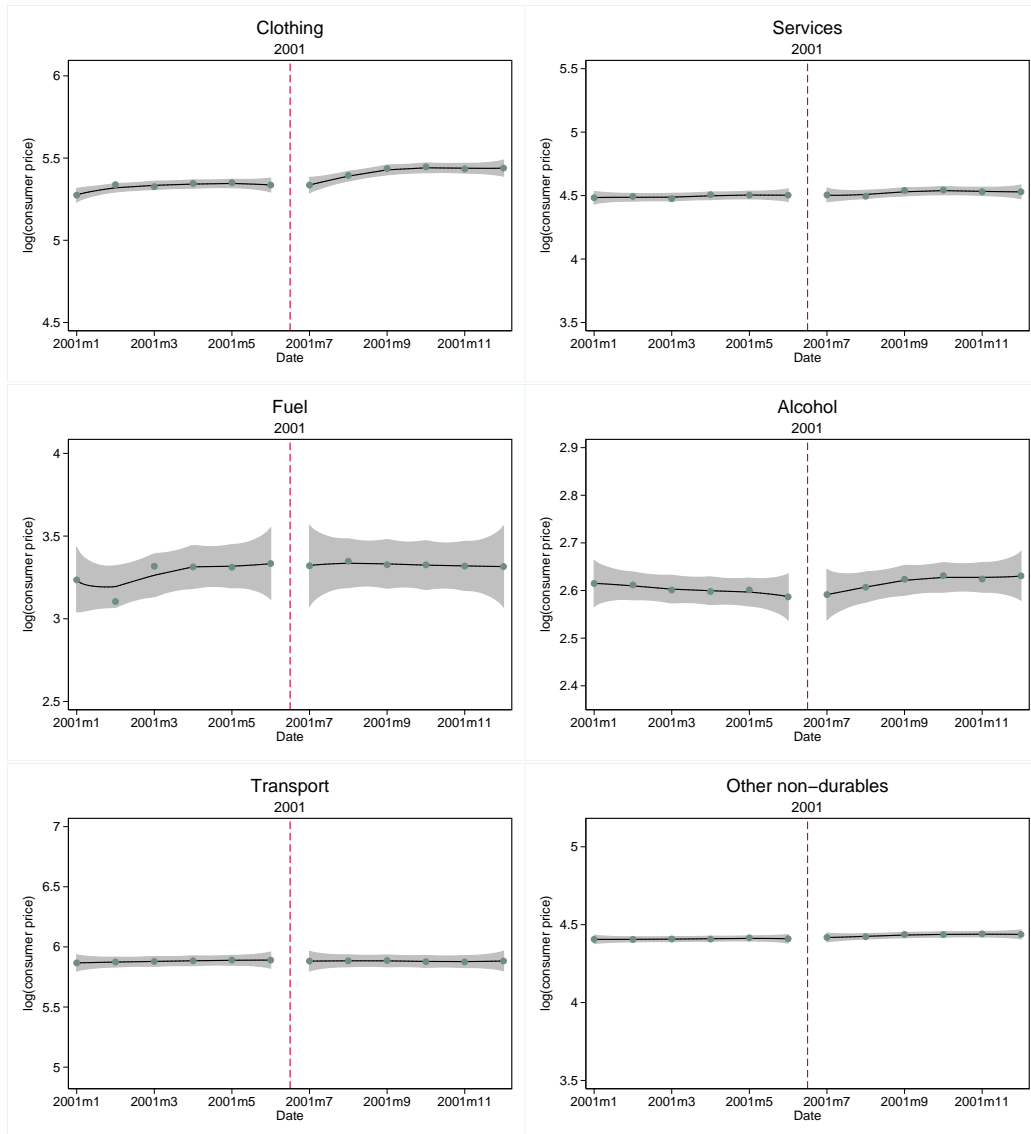
kernel applied to each side of the reform date: While the regression lines better illustrate the trends in the data and the size of the jumps at the reform date, the unrestricted means indicate the underlying noise in the data. The figure shows evidence of a sharp decline in food prices at the time of the reform, suggesting that VAT are heavily shifted to consumer prices. In contrast, there is no evidence of changes in food prices prior to the VAT reform. This suggests that firms do not change food prices in anticipation of the VAT reform. Additionally, food prices barely move in the months following the reform, suggesting that firms respond swiftly to the change in VAT.

In Figure 1.3, we present unrestricted and estimated monthly means of consumer prices for other goods. As in Figure 1.2, the scale of the y-axis is set equal to $\pm .5$ standard deviation of the respective variable. By standardizing the y-axes in this way, it becomes easier to compare the trends in the data and the sizes of the jumps at the cut-off dates across the graphs. In Figure 1.3, there is no sign of discontinuous changes in consumer prices around the reform date. This finding suggests the VAT reform did not affect the pricing of other goods that were not directly affected by the policy change.

Our RD design assumes that consumer prices would have evolved smoothly around the reform date in the absence of the policy change. This continuity condition implies that other observable determinants of consumer prices should have the same distribution just before and just after the reform. For simplicity, we consider a scalar representation of the observable determinants, given by the predictions from a regression of food prices on a rich set of covariates.⁷ The covariates are jointly predictive of food prices (F-statistic of 16.04). Figure 1.4 displays the predicted price in each month, showing no evidence of changes in observables around the time of the reform. Indeed, when looking at each covariate separately, we find no evidence of

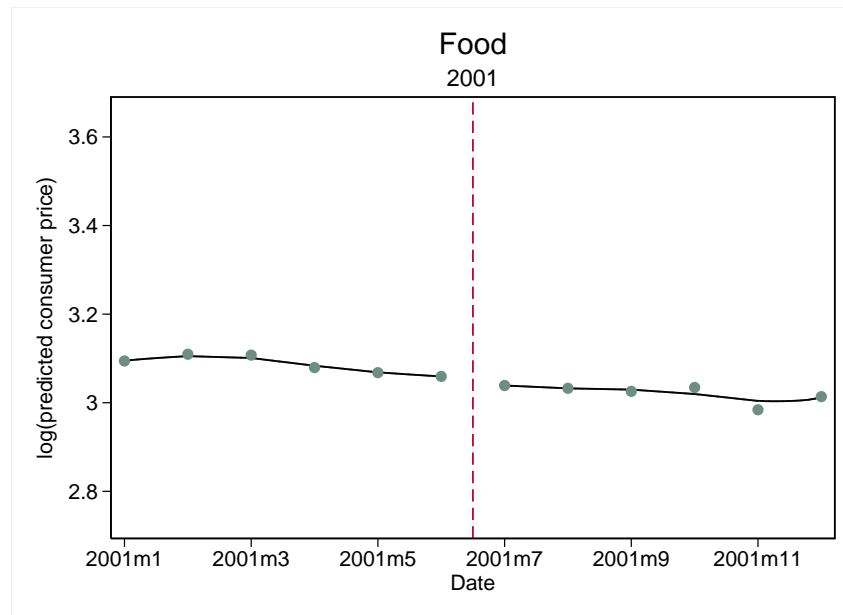
⁷We control for oil prices to proxy for energy and transportation prices. Since the price of imported food is likely to depend on the exchange rate, we control for the import weighted exchange rate as well as the exchange rates of Norway's key trading partners. All covariates are included with a one week and four weeks lag.

Figure 1.3: Evolution of consumer prices on non-food items over time



Notes: Each observation is the average consumer price for the specific good reported on the 15th each month. The dashed vertical lines denote the reform date. The solid lines are from a local linear regression with triangular weights on monthly consumer price data. The shaded area shows the 95 % confidence bands. The shaded area shows the 95 % confidence bands. The y-axes are scaled to $\pm .5$ st.dev. of the mean consumer price.

Figure 1.4: Evolution of predicted consumer price on food over time



Notes: Predicted food consumer prices is given by a regression of food consumer prices on oil prices (brent), euro/NOK, SEK/NOK, GDP/NOK and an import weighted exchange rate. All covariates are included with both one week and four weeks lag. Each observation is the average predicted consumer price for food reported on the 15th each month. The dashed vertical line denote the reform date. The solid lines are from a local linear regression with triangular weights on the predicted food consumer prices. The shaded area shows the 95 % confidence bands. The y-axes are scaled to $\pm .5$ st.dev. of the mean predicted consumer price.

any changes around the time of the reform. Nor do the estimates of λ change appreciatively if we add the covariates to equation (1.6).

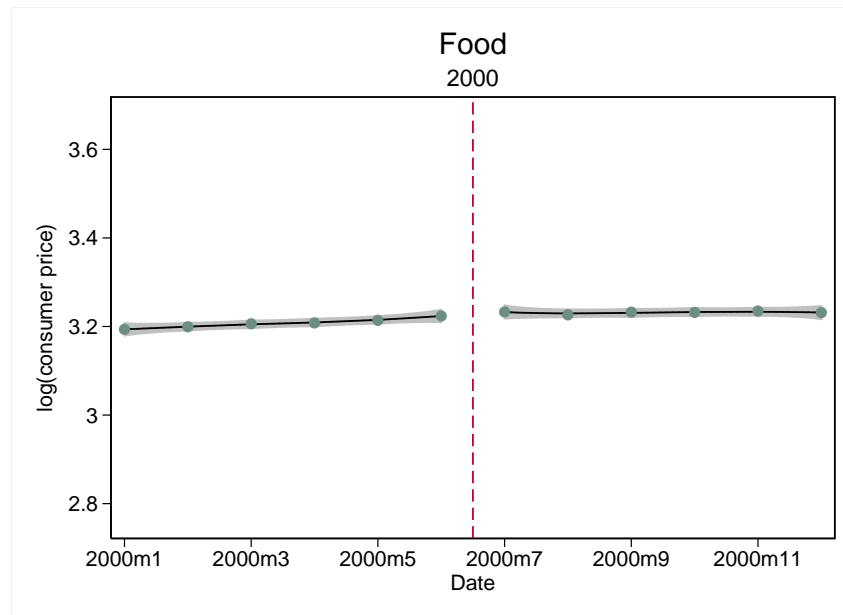
Another threat to our identification is that seasonality or month effects could generate discontinuous changes in food prices around the time of the reform. Although Figure 1.3 shows no sign of changes in food prices in the months before and after the reform, we cannot rule out that there would have been a discontinuous change in July 2001 in the absence of the policy change. To investigate the possibility of a July-specific month effect, Figure (1.5) shows the unrestricted and the estimated monthly means of consumer prices for food items during 2000. There is no evidence of any changes in food prices around July 1 of 2000; this finding is reassuring because there were no reform of the VAT system during this period.

1.3.3 Regression estimates

Having shown the raw patterns on the variables of interests around the reform date we now turn to regression-based estimates. The first column of Table 1.4 reports results from the FD model, comparing the consumer prices in June and July 2001. The first row shows the point estimate and standard error of the impact of the VAT reform on consumer prices for food. The point estimate suggests the reform reduced food prices by 11.2 percent. By way of comparison, full shifting would imply a reduction in food prices of 10.2 percent. This suggests that VAT on food items are completely shifted – or even slightly overshifted – to consumer prices. Indeed, the FD estimate is sufficiently precise to conclude that producers bear little, if anything, of the tax burden. Consistent with Figure 1.3, the other rows of Table 1.4 suggest the pricing of non-food items are not materially affected by the VAT reform.

The first column of Table 1.5 reports the RD estimates with 2 months of bandwidth on each side of the

Figure 1.5: Food consumer prices as function of month, 2000



Notes: Each observation is the average consumer price for food reported on the 15th each month. The dashed vertical line denote July 1st, 2000. The solid lines are from a local linear regression with triangular weights on monthly consumer price data. The shaded area shows the 95 % confidence bands. The shaded area shows the 95 % confidence bands. The y-axes are scaled to $\pm .5$ st.dev. of the mean consumer price.

reform date. The key difference between the FD model and the RD model is their assumptions regarding how the prices would have changed over time in the absence of the reform. The FD specification takes the consumer price on a good in June 2001 as a counterfactual for the price on the same good in July 2001. If there were secular changes in prices over this time period, the FD model would produce biased estimates of the effect of the VAT reform, because the price in June 2001 would be an inappropriate counterfactual for the price in July 2001. In this type of “smoothly contaminated” experiment, the RD specification uses the observed trends in prices on each side of the reform date to construct an appropriate counterfactual. As is evident from the first column of Table 2, the RD estimates are very similar to the FD estimates. The point estimate reported in the first row suggests the reform reduced food prices by 11.7 percentage points, which supports the conclusion that consumers bear the burden of VAT on food items. Additionally, the RD estimates confirm the finding of little if any spillover effects of the VAT reform to the consumer prices on non-food items.

1.3.4 Specification checks

To increase the confidence in our identification strategy, we now show that our regression estimates are robust to several specification checks.

We begin by adding a full set of item fixed effects to the regression models. The second column of Tables 1.4 and 1.5 report the results. We find that the estimates change little when including the fixed effects, suggesting the estimated reform effects are not driven by changes in the composition of commodities over time. However, including the fixed effects reduces the residual variance and is thus a useful way to gain precision. Indeed, when including commodity fixed effects we can rule out undershifting at conventional

Table 1.4: First difference results

Commodity:	(1)	(2)	(3)	Obs
Food	-0.112*** (0.0108)	-0.105*** (0.00360)	-0.107*** (0.00490)	34286
Clothing	-0.00735 (0.0328)	-0.0120 (0.0134)	0.00918 (0.0196)	7492
Services	-0.000980 (0.0392)	0.00275 (0.0119)	0.00848 (0.0172)	5225
Fuel	0.00801 (0.0220)	-0.00465 (0.0290)	0.0192 (0.0400)	300
Alcohol	0.00947 (0.0346)	0.00592 (0.0197)	-0.00199 (0.0277)	864
Transport	-0.0212 (0.0513)	-0.00663 (0.0110)	-0.0102 (0.0159)	6421
Other	0.00135 (0.0202)	0.00154 (0.00587)	0.00208 (0.00859)	21606
Item F.E.	No	Yes	Yes	
Month effect	No	No	Yes	

Notes: All coefficients are estimated using a FD model that compares log consumer prices in June 2001 with log consumer prices in July 2001. The first column report the results with no controls, the second column includes item fixed effects, and the third column also control for possible month effects using a DiD strategy. Standard errors are reported in parenthesis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

levels of significance.

Next, we estimate a difference-in-differences (DiD) specification of both the FD model and the RD model. The main motivation for this robustness check is that seasonality or month effects could generate discontinuous changes in consumer prices. The DiD specification exploits the fact that there were no change in the VAT rates in 2000: significant changes in the food prices in July 2000 would therefore be unrelated to the VAT system and should instead capture month effects. A standard DiD estimate is obtained by estimating the difference in consumer prices in July versus June for the year 2000, and subtract it from the FD estimate of the VAT reform. The DiD specification of the RD model is implemented in a similar way, by separately estimating equation (1.6) using data from 2000 or 2001. The third column of Tables 1.4 and 1.5 report estimates from the DiD specification of the FD model and the RD model, respectively. The results suggest that month effects do not confound the conclusions drawn about the passthrough of the VAT reform.

Lastly, we examine whether prices change in anticipation of the VAT reform. The change in tax rates was announced in December 2000, and it is conceivable that firms or consumers adjust their behavior prior to the reform date. However, the graphical evidence presented in Figure 1.2 showed no sign of changes in food prices outside the reform window. Further evidence against anticipation effects is provided by splitting the set of food items into fresh and storable food. The idea is that any anticipation effect should be stronger for storable food than for fresh food, and as a result, put downward pressure on the estimated passthrough of the VAT reform for storable food. However, when estimating the FD and RD models separately for storable

Table 1.5: Regression discontinuity results

Commodity:	(1)	(2)	(3)	Obs
Food	-0.117*** (0.0105)	-0.113*** (0.00366)	-0.104*** (0.00639)	93157
Clothing	-0.0101 (0.0318)	-0.0181 (0.0131)	0.0527 (0.0243)	21120
Services	-0.0109 (0.0367)	-0.00206 (0.0114)	0.00420 (0.0260)	14968
Fuel	-0.0306 (0.169)	-0.0287 (0.0314)	-0.0210 (0.0566)	917
Alcohol	0.00289 (0.0354)	0.00289 (0.0208)	0.0149 (0.0369)	2294
Transport	-0.00986 (0.0518)	-0.00986 (0.0113)	-0.00516 (0.0205)	17297
Other	0.00372 (0.0198)	0.00175 (0.00600)	0.00956 (0.0110)	60816
Commodity	No	Yes	Yes	
Month effect	No	No	Yes	

Notes: All coefficients are estimated using a RD model with log consumer prices as dependent variable, triangular weights and two months bandwidth. The first column report the results with no controls, the second column includes item fixed effects, and the third column also control for possible month effects using a DiD strategy. Standard errors are reported in parenthesis. *p < 0.1, **p < 0.5, ***p < 0.01.

and fresh food, we find very similar reform effects. The FD estimates of the reform effect is -0.085 (s.e.=0.0086) for storable goods, and -0.125 (s.e.=0.0178) for fresh goods. By comparison, the RD model produces estimates of the reform effect of -0.095 (s.e.=0.0168) for storable goods and -0.110 (s.e.=0.0338) for fresh goods.

1.4 Distributional effects of VAT reform

The RD estimates showed that the gains from the VAT reform ultimately fell on consumers rather than producers. In this section, we investigate how the passthrough to consumer prices affected the welfare of poor and rich households.

1.4.1 Model and estimates of demand system

Demand system. To study the welfare effects of the VAT reform, we apply the Almost Ideal (AI) demand system first proposed by Deaton and Muellbauer (1980). In the AI demand system, preferences belong to the Price-Independent Generalized Logarithmic (PIGLOG) class (Muellbauer, 1976) and they are defined by the expenditure function

$$\log c(u, \mathbf{p}) = (1 - u) \log a(\mathbf{p}) + u \log b(\mathbf{p}) \quad (1.7)$$

where u is the indirect utility and \mathbf{p} is a vector of prices of n goods. The functions $a(\mathbf{p})$ and $b(\mathbf{p})$ are specified by the following functional forms:

$$\log a(\mathbf{p}) = \alpha_0 + \sum_{i=1}^n \alpha_i \log p_i + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \gamma_{ij}^* \log p_i \log p_j, \quad (1.8)$$

and

$$\log b(\mathbf{p}) = \log a(\mathbf{p}) + \prod_{i=1}^n p_i^{\beta_i}. \quad (1.9)$$

In line with Browning and Meghir (1991) and Blundell, Pashardes, and Weber (1993), we use a two-stage budgeting framework. Preferences are characterized such that, in each period t , household h makes decisions on how much to consume of a set of non-durable commodities, conditional on various household characteristics and also conditional on the consumption levels of a second group of other (possibly less flexible) demands. The commodities we model directly (\mathbf{q}) are food, clothing, services, household fuel, alcohol, transport, and other non-durable goods. The second group contains housing, some durables, and labor-market decisions which, together with household characteristics, is represented by \mathbf{z} . Household utility is defined over \mathbf{q}_t^h for household h in period t conditional on the set of demographics and other conditioning variables \mathbf{z}_t^h .

The first stage of the budgeting framework is to allocate expenditures to commodities \mathbf{q}_t^h , denoted by m_t^h . The second stage of the budgeting framework is described by equations (1.8)-(1.10). Inserting for (1.8) and (1.9) in (1.7) and applying Roy's identity gives the budget shares

$$w_{it}^h = \alpha_{it}^h + \sum_j \gamma_{ij} \ln p_{jt} + \beta_{it}^h \ln \left[\frac{m_t^h}{a(\mathbf{p})} \right] \quad (1.10)$$

where w_{it}^h is household h 's budget share of good i , and p_{jt} is the price of good j at time t . The term $[m_t^h/a(\mathbf{p})]$ represents relative income with $a(\mathbf{p})$ being a price index. Household preferences are incorporated by allowing the constant α_{it}^h to depend on household characteristics, z_{kt}^h ,

$$\alpha_{it}^h = \alpha_i + \sum_k \alpha_{ik} z_{kt}^h + \sum_k \delta_k T_{kt},$$

in which we have also added full set of indicator variables for year and season T_{kt} .

Both the indirect utility function and the demand functions for each good that arise from equation (1.7) are linear in the log of total expenditure. Figure (1.1) in Section 2 examined this assumption for our main commodity of interest, food items, by looking at the Engel curve. This figure provided a nonparametric description of the Engel curve and shows that the linear model seems to be a reasonable approximation for the food share curve.⁸

Estimation procedure. To consistently estimate (γ_{ij}, β_i) for every commodity i , we use the two-step estimation method of Browning and Meghir (1991) and Blundell, Pashardes, and Weber (1993). This estimation method incorporates a set of theoretical within-equation and cross-equation restrictions. Furthermore, it accounts for the endogeneity of $\ln m_t^h$ in the budget share equations.

⁸This is consistent with what Banks, Blundell, and Lewbel (1997) find using British household data.

The first step imposes the within-equation restrictions of adding-up and (zero-degree) homogeneity on (1.10) by expressing all prices relative to the price of “other” goods together with excluding this equation from the system. Each equation is estimated separately, allowing for heteroscedasticity in the error terms. We use GMM to obtain unrestricted consistent estimates for each equation. Additionally, an iterative method is applied where one take advantage of the conditional linearity of equation (1.10) given $a(\mathbf{p})$. That is, given $a(\mathbf{p})$, the system is linear in parameters, and this suggests a natural iterative procedure conditioning on an updates $a(\mathbf{p})$ at each iteration.⁹

The second step imposes the cross-equation restriction of symmetry. Let ϕ (ϕ^*) denote the vector of unrestricted (restricted) parameters obtained in the step outlined above. The cross-equation restrictions on ϕ can then be expressed as

$$\phi = \mathbf{K}\phi^*, \quad (1.11)$$

where \mathbf{K} is a matrix of rank $l - m(m - 1)/2$ and l is the number of unrestricted parameters in the demand equation system. To impose these restrictions the MCS method chooses an estimator $\hat{\phi}^*$ so as to minimize the quadratic form

$$\hat{\phi}^* = \arg \min \left[\hat{\phi} - \mathbf{K}\phi^* \right]' \boldsymbol{\Sigma}_\phi^{-1} \left[\hat{\phi} - \mathbf{K}\phi^* \right] \quad (1.12)$$

where $\hat{\phi}$ is the vector of unrestricted parameter estimates and $\boldsymbol{\Sigma}_\phi$ is its estimated covariance matrix. An estimate of the covariance matrix of the symmetry constrained estimator is $(\mathbf{K}'\boldsymbol{\Sigma}_\phi^{-1}\mathbf{K})^{-1}$.

Parameter estimates and elasticities. We now turn to the estimated parameters and implied elasticities of the individual household allocations. The price and income coefficients that correspond to the γ_{ij} and β_i parameters in equation (1.10) are given in Table 1.6. From the parameter estimates we see that food and fuel are necessities, while clothing, services and transport are luxury goods.

To interpret the parameters, it is useful to consider the implied elasticities. The budget elasticity at reference price is defined as

$$\varepsilon_i^h = \frac{\beta_i}{w_i^h} + 1.$$

The budget elasticities will vary with family composition since the predicted expenditure share w_i^h varies across households. The uncompensated demand elasticity of good i w.r.t. the price of good j at reference prices is given by

$$\varepsilon_{ij}^u = \frac{1}{w_i^h} \left[\gamma_{ij} - \beta_i \left(\alpha_j + \sum_{k=1}^N \gamma_{jk} \ln p_k \right) \right] - \delta_{ij}$$

where δ_{ij} is the Kronecker delta. Again, we see that the elasticities vary across households due to different budget shares. The compensated price elasticity is

$$\varepsilon_{ij}^c = \varepsilon_{ij}^u + \varepsilon_i w_j^h,$$

where the compensated price elasticity allows the consumers to revise their expenditure decision made in stage one of the budgeting framework when the price of good j changes. The elasticities are reported in table 1.7. These elasticities are calculated for each household individually, and then a weighted average is constructed, with the weights being equal to the household’s share of total sample expenditure of the relevant good. As expected, the uncompensated and compensated own price elasticities are negative for all goods. In terms of magnitudes, the elasticity estimates are in line with the findings of Banks, Blundell, and Lewbel

⁹As a first approximation to $a(\mathbf{p})$, we compute household-specific Stone price indices.

Table 1.6: Estimated parameters from the demand system

	Share equations					
	Food	Clothing	Services	Fuel	Alcohol	Transport
Expenditure	-0.0727*** (0.0138)	0.0453*** (0.0094)	0.0236*** (0.0070)	-0.0793*** (0.0079)	0.0023 (0.0022)	0.0589*** (0.0117)
Price food	-0.0750 (0.0465)	0.0629* (0.0285)	0.0041 (0.0196)	-0.0828*** (0.0170)	-0.0032 (0.0082)	0.0719* (0.0339)
Price clothing	0.0629* (0.0285)	-0.0558* (0.0270)	0.0060 (0.0130)	0.0543*** (0.0112)	0.0115* (0.0055)	-0.0828** (0.0255)
Price services	0.0041 (0.0196)	0.0060 (0.0130)	0.0002 (0.0151)	0.0226** (0.0086)	-0.0024 (0.0039)	-0.0245 (0.0228)
Price fuel	-0.0828*** (0.0170)	0.0543*** (0.0112)	0.0226** (0.0086)	-0.0938*** (0.0095)	0.0031 (0.0026)	0.0707*** (0.0140)
Price alcohol	-0.0032 (0.0082)	0.0115* (0.0055)	-0.0024 (0.0039)	0.0031 (0.0026)	-0.0055 (0.0057)	-0.0005 (0.0083)
Price transport	0.0719* (0.0339)	-0.0828** (0.0255)	-0.0245 (0.0228)	0.0707*** (0.0140)	-0.0005 (0.0083)	-0.0446 (0.0494)

Notes: γ -symmetry constrained Almost Ideal estimates. Standard errors are reported in parenthesis. Control variables include: year and seasonal dummies, age and gender of the head of household, number of cars in the household, working status of the wife, number of children age 0-7 years, 7-16 years and 16-20 years. Dummy for smoking, single parent and whether the head of household is retired. Geographically dummies for region and size of city/village. Total expenditures on non-durables are treated endogenous. Excluded instrument is household income. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

(1997)

1.4.2 Distributional effects

First order approximation. We begin with a first order approximation of the distributional effects, ignoring any behavioral responses to the reduction in prices. To this end, we multiple the pre-reform expenditure levels on the various goods with the RD estimates of the prices changes. After the VAT reform, households can buy the same bundle of goods at a lower price. Assuming no behavioral responses and ignoring cross-price effects, the VAT reform amounts to a cash transfer equal to 11.7 % of a household's expenditure on food. The first three columns in Panel A of Table 1.8 illustrates the variation in the size of this transfer across households; these columns report the 25th percentile, the median, and the 75th percentile of the distribution of transfer amounts. We see that the size of the transfer is NOK 3222 at the 25th percentile, whereas the transfer at the 75th percentile is about twice as large. By comparison, the median transfer is NOK 4864. In the last three columns, we also take into account the estimated price changes on non-food items and the households' pre-reform expenditure levels on these goods. As expected, we find that incorporating these negative cross-price effects increases the value of the transfer.

Allowing for behavioral responses. There is an obvious attraction to simply using information on observed expenditure patterns to assess the welfare implications of the VAT reform. No response parameters are required, and therefore the analysis is not subject to estimation error in own or cross-price demand elasticities. However, the VAT reform generated substantial rather than marginal changes in food prices. In such cases, substitution effects can be non-trivial, as consumers substitute towards relatively cheaper goods. The first order approximations ignore these effects, and therefore, can be seriously biased (see e.g. Banks,

Table 1.7: Implied elasticities from the parameters of the demand system

A: Budget elasticities:							
	Commodity						
	Food	Clothing	Services	Fuel	Alcohol	Transport	Other
	0.694	1.397	1.139	-0.165	1.134	1.355	1.088
B: Uncompensated cross price elasticities:							
	Commodity						
Price	Food	Clothing	Services	Fuel	Alcohol	Transport	Other
Food	-1.164	0.254	0.043	-0.247	-0.012	0.291	0.142
Clothing	0.354	-1.474	0.019	0.343	0.100	-0.709	-0.030
Service	-0.045	0.041	-1.011	0.086	-0.015	-0.139	-0.057
Fuel	-0.638	0.754	0.430	-1.989	0.049	0.991	0.567
Alcohol	-0.255	0.684	-0.154	0.135	-1.321	-0.025	-0.199
Transport	0.258	-0.486	-0.178	0.308	-0.004	-1.255	0.003
Other	0.045	0.019	-0.031	0.075	-0.012	0.043	-1.225
C: Compensated cross price elasticities:							
	Commodity						
Price	Food	Clothing	Services	Fuel	Alcohol	Transport	Other
Food	-0.945	0.325	0.143	-0.192	-0.005	0.381	0.292
Clothing	0.671	-1.235	0.237	0.432	0.113	-0.513	0.295
Services	0.202	0.158	-0.729	0.159	-0.003	0.025	0.188
Fuel	-0.650	0.728	0.392	-1.973	0.048	0.950	0.505
Alcohol	0.032	0.799	0.035	0.213	-1.276	0.138	0.059
Transport	0.549	-0.353	0.030	0.393	0.008	-0.917	0.289
Other	0.281	0.129	0.123	0.143	-0.002	0.182	-0.854

Notes: Calculated elasticities using the γ -symmetry constrained Almost Ideal estimates.

Blundell, and Lewbel, 1996).

To allow for behavioral responses, we use the parameter estimates of the AI model to calculate the indirect utility of the households from equations (1.7) - (1.9). The transfer to a given household with characteristics z is measured as the compensating variation, given by the difference in the cost functions $c(\mathbf{p}^0, z, u^0) - c(\mathbf{p}^1, z, u^0)$, where the post-reform cost function is evaluated at the pre-reform indirect utility level. This welfare measure tells us the maximum amount of income a household is willing to pay for the VAT reform. The first three columns of Panel B in Table 1.8 illustrates the variation in the compensating variation across households when we abstract from the cross-price effects. Because a household may adjust its spending pattern in response to changes in prices, the willingness to pay exceeds 11.7 % of the household's expenditure on food prior to the reform. For example, the compensating variation amounts to NOK 7616 at the 25th percentile and NOK 14379 at the 75th percentile. As above, incorporating the negative cross-price increases further the willingness to pay for the VAT reform.

Distributional effects

We begin with a graphical depiction of the distributional effects of the VAT reform, before quantifying how its impact on inequality. Figure 1.6 plots the magnitude of the compensating variation against total expenditure. As rich households consume more food, the willingness to pay increases with total expenditure. Figure 1.7 complements by showing the relative size of the compensating variation. This figure reveals that richer

Table 1.8: Size of transfers

	percentile			percentile		
	25th	50th	75th	25th	50th	75th
<i>Panel A: First-order approximation</i>						
Transfer in NOK	3222	4864	6836	4151	5970	8066
<i>Panel B: Behavioral response</i>						
Compensating variation in NOK	7616	10881	14379	8914	13085	17715
Indirect price responses	No			Yes		

Notes: Column 1-3 show size of transfer when only allowing for direct price response to the VAT reform. Column 4-6 show size of transfer allowing for both direct and indirect price responses to the VAT reform.

Table 1.9: Percentage change in Gini coefficient

<i>Panel A: First-order approximation</i>			
Δ Gini coefficient		-0.48 %	-0.58 %
<i>Panel B: Behavioral response</i>			
Δ Gini coefficient		-0.90 %	-0.98 %
Indirect price responses	No		Yes

Notes: Column 1 shows the percentage change in the GINI coefficient when only allowing for direct price response to the VAT reform. Column 2 shows the percentage change in the GINI coefficient allowing for both direct and indirect price responses to the VAT reform.


households are willing to pay a smaller fraction of their total expenditure for the VAT reform.

To summarize the impact of the VAT reform on inequality, we employ the much used Gini coefficient. In 2000, the distribution of household income in our sample gives a Gini coefficient of 0.210. To assess the impact of the VAT reform, we add the size of the transfer to each household and compute the Gini coefficient in this simulated distribution of household income. Panel A of Table 1.9 reports the change in the Gini coefficient when the size of the transfer is computed by the first order approximation. If we abstract from cross-price effects, the Gini coefficient is reduced by 0.48 % when we include these transfers. Allowing for the cross-price effects, reduces the Gini coefficient further. Panel B of Table 1.9 reports the change in the Gini coefficient when the size of the transfer is computed by compensating variations. These transfers are substantially larger in absolute amounts and they have a larger impact on the distribution of household income. Column 1 abstracts from cross-price effects and shows that the reduction in the Gini coefficient is equal to 0.9 %. In column 2, we allow for cross-price effects and the result is that the Gini coefficient is reduced further to 0.98 %. Put into perspective, this reduction in the Gini coefficient corresponds to introducing a 0.98 % proportional tax on earnings and then redistributing the derived tax revenue as equal sized amounts to the individuals (Aaberge, 1997).

1.5 Conclusion

Much of the controversy surrounding recent policy proposals to broaden the base for value added taxes (VAT) revolves around who ultimately bears the burden of these taxes. The typical assumption is that consumer prices fully reflect taxes, so that the only important empirical question is how the tax induced price changes affect members of different income groups. For example, the Mirrless Review assumes the incidence is fully

Figure 1.6: Compensating variations. NOK per year.




Notes: Compensating variations when allowing for both direct and indirect price responses to the VAT reform. In NOK per year. Slope coefficient is 0.056 (s.e.=0.00012)

on consumer prices in their proposal to broaden the base for VAT by removing the zero rating for food. However, the evidence base is scarce, and as critics of such policy changes point out (see e.g. Atkinson, 2013), market imperfections could generate both over and undershifting of VAT to consumer prices.

In this paper, we examined the incidence and distributional effects of VAT in a setting with plausibly exogenous variation in tax rates. The context of our study was a sharp change in the VAT policy on food items in Norway. Using a regression discontinuity design, we examined the direct impact of the policy change on the consumer prices of food items as well as any cross-price effects on other goods. Our estimates suggested that taxes levied on food items are completely shifted to consumer prices, whereas the pricing of other goods is not materially affected. To understand the distributional effects of the VAT reform, we used expenditure data and estimated the compensating variation of the tax induced price changes. We found that lowering the VAT on food attenuates inequality in consumer welfare, in part because households adjust their spending patterns in response to the prices changes. By comparison, the usual first order approximation of the distributional effects, which ignores behavioural responses, seriously understates the redistributive nature of the VAT reform.

Figure 1.7: Relative compensating variations.



Notes: Allowing for both direct and indirect price responses to the VAT reform. Relative compensating variation is defined as compensating variation/total expenditure.

Chapter 2

The Skill Complementarity of Broadband Internet¹

2.1 Introduction

Economists and policymakers are keenly interested in understanding the productivity and labor market effects of the advancements in information and communication technology (ICT). Many have argued that these technological changes were behind the resurgence in U.S. productivity growth since the late 1990s, and that investments in ICT are important in explaining labor productivity patterns across multiple industries and countries.² Recently, policymakers have pointed to broadband internet as a key productivity enhancing factor, calling for public funding to roll out broadband infrastructure.³ While government agencies are projecting broadband penetration rates to be important for productivity and job creation, there is little scientific evidence to substantiate these claims.⁴

Estimating how adoption of new technology, such as broadband internet, affects productivity has proven difficult for several reasons. It is often difficult to access data on technology adoption which can be linked with firm-level information on value added and factor inputs. Another key challenge is the likelihood that some determinants of production are unobserved to the econometrician but observed by the firm; if adoption of new technology depends on these determinants, then OLS estimates of production functions will be biased. On top of this, the use of inputs such as capital and labor could also be correlated with technology adoption and unobserved productivity, and therefore create bias in OLS estimates. These empirical challenges have meant that existing research has largely focused on demonstrating positive associations of ICT with productivity.

¹Joint work with Anders Akerman (Stockholm University) and Magne Mogstad (University of Chicago).

²The evidence is reviewed in Draca, Sadun, and Van Reenen (2007), Oliner, Sichel, and Stiroh (2007), Jorgenson, Ho, and Stiroh (2008), and Syverson (2011). See also Bloom, Sadun, and Van Reenen (2012).

³President Obama's campaign in 2008 promised that "every American should have the highest form of broadband access". Later, he called on Congress to approve funding for rolling out broadband to unserved and underserved areas, arguing that increased broadband spending will keep America competitive while also creating new jobs. In the UK, former Prime Minister Gordon Brown compared the expansion of broadband infrastructure to "the roads and the bridges and the railways that were built in previous times to stimulate the economy".

⁴In 2008, the U.S. Commissioner of Federal Communications Commission stated that "Affordable broadband would quickly add \$500 billion to the U.S. Economy and create 1.2 million jobs". Projections from the U.S. Bureau of Economic Analysis suggest that for every \$1 invested in broadband, the economy benefits nearly \$3. A notable example of research on this topic is Czernich, Falck, Kretschmer, and Woessmann (2011), who find positive association between broadband penetration and economic growth across OECD countries over time. Another example is Forman, Goldfarb, and Greenstein (2012), showing that internet investments correlate with wage and employment growth in some but not all U.S. counties.

Draca, Sadun, and Van Reenen (2007), in their Handbook of Information and Communication Technologies chapter, review the literature carefully and conclude that “none of the literature has produced convincing evidence of a causal impact of ICT on productivity, for example by analyzing a natural experiment”.

In this paper, we examine how the adoption of new technology in firms shifts their production technology and changes the productivity and labor outcomes of different types of workers. Our context is the adoption of broadband internet in Norwegian firms over the period 2001-2007.⁵ Norway is a small open economy with segmented local labor markets. Our analysis employs several data sources that we can link through unique firm and individual identifiers. Using firm-level information on value added, capital, labor by skill level and broadband adoption, we estimate production functions where firms can change their technology by adopting broadband internet.⁶ We complement these production function estimates with wage regressions which include interactions between skill levels and broadband internet.

As a source of exogenous variation in broadband adoption, we follow Bhuller, Havnes, Leuven, and Mogstad (2013) in exploiting a public program aimed at ensuring broadband access at a reasonable price to all households throughout the country.⁷ Because of limited funding, access to broadband was progressively rolled out, so that the necessary infrastructure (access points) was established in different municipalities at different times. We use this spatial and temporal variation in the availability of broadband across municipalities as an instrumental variable for adoption of broadband in firms. Our identification strategy – which controls for municipality (time and industry) fixed effects – is motivated by two features of the broadband program. First, most of the supply and demand factors tend to vary little over time. Second, the timing of the roll-out is unlikely to co-vary with the key correlates of productivity and labor outcomes. We demonstrate that the data is consistent with these program features, and further challenge the validity of the instrument by probing the stability of the estimates to alternative specifications, finding little cause for worry. Following Levinsohn and Petrin (2003), we also make sure that the results are robust to using intermediate inputs to proxy for unobserved productivity in the production function.

Figure 2.1 previews the main results of our analysis. This figure shows estimates of labor productivity and skill premium before and after the expansion of broadband internet, providing a visual representation of the reduced form of our econometric models. Time zero represents the year with the strongest growth in broadband coverage rate in a given municipality, after taking out municipality, industry and year fixed effects. In each time period, we estimate Cobb-Douglas production functions and wage regressions while controlling for municipality, industry and year fixed effects. The graphs suggest strong skill complementarity of broadband internet. The first (second) graph shows that the sharp rise in coverage rates from time -1 to time zero is associated with a substantial increase (decrease) in the output elasticity of (un)skilled labor. The third graph shows how relative wages of skilled workers increases after this rise in coverage rates.

Consistent with this reduced form evidence, the structural production function estimates show that broadband adoption in firms favors skilled labor by increasing its relative productivity. The estimated increase in productivity of skilled labor is especially large for workers with college degree in high return fields such as science, technology, engineering, and business. By comparison, broadband internet is a substitute for workers without high school diploma, lowering their marginal productivity. In line with the estimated changes in

⁵Throughout the paper, we focus on broadband internet, defined as internet connections with download speeds that exceed 256 kbit/s. Before the expansion of broadband internet, all firms with a telephone connection would have dial-up access to internet, but limited to a bitrate of less than 56 kbit/s. Broadband internet facilitated internet use without excessive waiting times.

⁶Throughout the paper, we follow the literature by referring to education and skills interchangeably; thus skilled refers to highly educated and unskilled refers to those with lower levels of education.

⁷Bhuller, Havnes, Leuven, and Mogstad (2013) use the roll-out of broadband internet to study how internet use affects sex crimes.

labor productivity, the wage regressions show the expansion of broadband internet improves (worsens) the labor outcomes of (un)skilled workers. The gains are largest for college graduates in high return fields.

Our findings are consistent with the widespread view that ICT is complementary with human capital. Technological changes that reduce quality-adjusted ICT prices – such as the arrival of broadband internet – should therefore increase skill demand and returns to skill. While the labor market consequences of investments in computers and R&D have been studied extensively,⁸ we provide novel evidence on the skill bias of broadband adoption. We also show an important link between computerization and broadband adoption in firms: When broadband internet becomes available, it is not randomly adopted; instead, it is more quickly adopted in firms in which complementary factors are abundant, including computers and skilled workers. These findings conform to the predictions of a model of endogenous technology adoption where firms' choices reflect principles of comparative advantage (see e.g. Beaudry and Green, 2003, 2005; Beaudry, Doms, and Lewis, 2010). This pattern of broadband adoption illustrates that the estimated effects pertain to a subgroup of early adopters of broadband internet, and that we need to be cautious in extrapolating these local estimates to the longer-run impact of broadband adoption for the population of firms at large.

We further use the rich Norwegian data to explore possible mechanisms behind the skill bias of broadband adoption. We find two pieces of evidence against the hypothesis that our findings reflect changes in the demand for goods due to the expansion of broadband internet. First, we find no impact of broadband adoption on online sales or procurement. Second, when we look separately at firms in the tradeable sector where demand is given by the world market, we still find that broadband adoption in firms favors skilled labor by increasing its relative productivity. Our data is also at odds with broadband adoption coinciding with general technical upgrading in firms. Another possible mechanism receiving little support in data is that our findings are driven by changes in labor productivity and wages in firms directly affected by the expansion of broadband access, such as telecom firms or IT consultancy companies.

Instead, our findings are consistent with the task approach to skill biased technological change. Autor, Levy, and Murnane (2003) argue that ICT substitutes for workers in performing routine tasks – more amenable to automatization – and complements workers in executing problem-solving, complex communication, and information-intensive tasks (ofted called “nonroutine abstract tasks”).⁹ We first show that workers of different educational background indeed cluster disproportionately in occupations that require different tasks. In particular, college graduates tend to move into occupations that are pervasive at nonroutine abstract tasks, especially those with degrees in high return fields. By comparison, low educated workers are over-represented in occupations that involve relatively little nonroutine abstract tasks, but require more routine or manual tasks. Next, we estimate wage regressions which include interactions between broadband internet and the tasks performed in jobs rather than (or in addition to) the educational credentials of workers performing those jobs. These estimates suggest an important channel behind the skill bias of broadband internet is that it complements non-routine abstract tasks but substitutes for routine tasks whilst not affecting manual tasks.

Taken together, our findings have important implications for the debate about the role of government policies in encouraging productivity growth.¹⁰ Our estimates suggest that policy increasing the broadband

⁸See e.g. Krueger (1993), Berman, Bound, and Griliches (1994), DiNardo and Pischke (1997), Autor, Katz, and Krueger (1998), Machin and Van Reenen (1998), Beaudry and Green (2003; 2005) and Beaudry, Doms, and Lewis (2010). Acemoglu (2003), Bond and Van Reenen (2007) and Goldin and Katz (2007) provide recent reviews of the extensive literature on technology-skill complementarity.

⁹See also Autor, Katz and Kearney (2006, 2008), Goos and Manning (2007), Black and Spitz-Oener (2010), Firpo, Fortin, and Lemieux (2011), Acemoglu and Autor (2011), Autor and Dorn (2013), and Michaels, Natraj, and Van Reenen (forthcoming). A related literature argues that ICT changes workplace organization and practices, by increasing skill requirements, worker autonomy and management's ability to monitor workers (see e.g. Caroli and Van Reenen, 2001; Bresnahan, Brynjolfsson, and Hitt, 2002; Brynjolfsson and Hitt, 2003; Bloom, Garicano, Sadun, and Van Reenen, 2009).

¹⁰Several studies have evaluated reforms with plausibly productivity enhancing effects (see the discussion in Syverson, 2011).

penetration rates could enhance firm productivity. A related issue is why policy changes, even if they encourage productivity, do not always happen. One explanation is that established interests earning rents in the unreformed environment could be able to stave off reform. Examining who wins, who loses, and by how much could inform about the nature of these barriers and how to design compensation schemes. Our study points to the skill bias of broadband induced shift in production technology as a barrier to government investment in broadband infrastructure.

The paper unfolds as follows. Section 2 presents a framework for estimation of technological change in production and discusses previous research. Section 3 describes our data, before Section 4 discusses the expansion of broadband internet. Section 5 presents our empirical strategy. Section 6 describes our main findings, discusses their economic significance, and reports results from a number of robustness checks. Section 7 explores possible mechanisms. The final section offers some concluding remarks.

2.2 Technological change in production

2.2.1 A framework for estimation

Economic theory views the production technology as a function describing how a collection of factor inputs can be transformed into output, and it defines a technological change as a shift in the production function, i.e., a change in output for given inputs. Consider a Cobb-Douglas production technology with exponents that potentially change with the adoption of broadband internet:

$$Y_{it} = e^{\alpha_0 + D_{it}\alpha_1} K_{it}^{\beta_{k0} + D_{it}\beta_{k1}} U_{it}^{\beta_{u0} + D_{it}\beta_{u1}} S_{it}^{\beta_{s0} + D_{it}\beta_{s1}} e^{\varepsilon_{it}}, \quad (2.1)$$

where Y_{it} represents output of firm i in period t , K_{it} , U_{it} and S_{it} are inputs of capital, unskilled, and skilled labor, respectively, $e^{\alpha_0 + D_{it}\alpha_1}$ is the average total factor productivity (TFP), and $e^{\varepsilon_{it}}$ is the time- and firm-specific deviation from average TFP. Equation (2.1) extends on the usual Cobb-Douglas technology through the indicator variable D_{it} , which is equal to one if firm i has broadband internet in period t (and zero otherwise). This simple extension allows for the possibility that broadband adoption is both factor-neutral ($\alpha_1 \neq 0$) and factor-biased ($\beta_{k1} \neq 0$, $\beta_{u1} \neq 0$ or $\beta_{s1} \neq 0$).

By estimating production functions like (2.1), we will explore several forms of factor-biased technological change that are discussed in the literature (see e.g. Berman, 2000). Broadband adoption is absolutely skill biased if the output elasticity of skilled workers increases,

$$\beta_{s1} > 0,$$

while broadband adoption is relatively skill biased if

$$\frac{\beta_{s1}}{\beta_{s0}} > \frac{\beta_{u1}}{\beta_{u0}},$$

and broadband adoption is strongly skill biased if

$$\beta_{s1} > 0 > \beta_{u1}.$$

One example is Olley and Pakes (1996), examining the impact of deregulation on productivity growth in the U.S. telecommunication industry. Another example is Pavcnik (2002), investigating the effects of liberalized trade on firm productivity in Chile.

A complementary approach to assess how broadband adoption shifts the production technology is to examine the labor market implications. Under standard conditions, relative skill bias implies that relative wages of skilled workers increase, holding the ratio of labor inputs fixed; or that the employment share of skilled workers increases, keeping relative wages constant. These two implications have been treated as alternative symptoms of skill-biased technological change in the literature. By comparison, strong skill bias implies that wages or employment rate of unskilled workers decreases with broadband adoption. By evaluating how the staged expansion of broadband internet affected hourly wages and employment rates in local labor markets, we will examine the implications of different forms of skill-biased technological change.

2.2.2 Empirical challenges and previous research

To estimate how broadband adoption shifts the production technology, it is useful to transform (2.1) to a linear production function:

$$y_{it} = \alpha_0 + \beta_{k0}k_{it} + \beta_{u0}u_{it} + \beta_{s0}s_{it} + D_{it}[\alpha_1 + \beta_{k1}k_{it} + \beta_{s1}s_{it} + \beta_{u1}u_{it}] + \varepsilon_{it}, \quad (2.2)$$

where lower-case letters refer to natural logarithms. While there exists a large literature on the identification and estimation of production functions,¹¹ there is little direct evidence on how adoption of new technology changes the productivity of different types of workers.

One reason is the difficulty in accessing data on technology adoption which can be linked with firm-level information on value added, capital, and labor by skill level. For instance, typical annual accounts data reports labor costs as average (or total) wages for all workers, rather than wages by the skill level of the workers (Van Beveren, 2012). Moreover, statistical agencies have only recently started to systematically collect ICT information at the firm level (Draca, Sadun, and Van Reenen, 2007). In most cases, this information relates to broad measures of expenditure or usage of ICT, rather than precise measures of specific technological changes.

Another reason is the threats to identification of production functions. A key concern is the likelihood that some determinants of production are unobserved to the econometrician but observed by the firm; if adoption of new technology depends on unobserved productivity, then OLS estimates of (2.2) will be biased. Several studies illustrate the difficulty in drawing credible inferences absent an appropriate instrumental variable. DiNardo and Pischke (1997) suggest that computer users possess unobserved skills which might have little to do with computers but which raise their productivity. Beaudry and Green (2003; 2005) and Beaudry, Doms, and Lewis (2010) show the speed and extent of computer adoption are likely endogenous and reflect principles of comparative advantage. Another generic concern in estimation of production functions is that the levels of capital and labor chosen could be correlated with unobserved factors of productivity. Olley and Pakes (1996) and Levinsohn and Petrin (2003) develop methods for estimating production functions that control for correlations between the use of factor inputs and unobserved productivity. However, most studies of technological change in production ignore these threats to identification and simply estimate a production function using OLS (Draca, Sadun, and Van Reenen, 2007). In particular, there have been few attempts to find instruments that affect the decision to adopt new technology.

Because we have data with firm-level information on value added, capital, labor by skill level and broadband adoption, we are able to estimate production functions like (2.2); as a source of exogenous variation in broadband adoption, we exploit that the necessary infrastructure was established in different areas at

¹¹For recent reviews, see Draca, Sadun, and Van Reenen (2007) and Van Beveren (2012).

different times; and following Levinsohn and Petrin (2003), we use intermediate inputs to proxy for unobserved productivity in the production function. This enables us to address the threats to identification and provide novel evidence on how broadband adoption in firms shifts the production technology and changes the productivity of different types of workers. The wage regressions which include interactions between skill levels and broadband internet allow us to examine the implications of skill-biased change without invoking the full set of assumptions behind the production function estimation.

2.3 Data

Our analysis uses several data sources, which we can link through unique identifiers for each firm, employee, and municipality. The coverage and reliability of Norwegian data are rated as exceptional in international quality assessments (see e.g. Atkinson, Rainwater, and Smeeding, 1995). Below we describe our data and sample selection, while details about the data sources and each of the variables are given in Appendix Table A1.

Firm and worker data

Our firm data comes from administrative registers, which are updated annually by Statistics Norway and verified by the Norwegian Tax Authority. The data comprises all non-financial joint-stock firms over the period 1999–2008.¹² It contains detailed information from the firm’s balance sheets on output (such as revenues) and inputs (such as capital, labor, intermediates) as well as 4-digit industry codes and geographical identifiers at the municipality level.

We merge the firm data set with a linked employer-employee registry that contains complete records of all firms and workers for the period 1999–2008. For every employee, we know his or her length and field of education, and annual labor income. In our baseline specification, we define an employee as skilled if he or she has completed high school, while employees with less schooling are defined as unskilled. In much of our analysis, we refine these often used proxies for skill levels: We divide skilled workers into medium skilled (high school graduates) and high skilled (post secondary education); and we interact post-secondary education with field of study.

Internet data

For the period 2001–2007, we have (i) data on broadband subscription for a stratified random sample of firms, and (ii) complete records of the fraction of households that have access to broadband internet in every municipality. As explained in detail below, we will use the former to measure broadband adoption in firms, while the latter will form the basis for our instrumental variable. Throughout this paper, broadband is defined as internet connections with download speed that exceeds 256 kbit/s.

Our data on broadband subscriptions of firms comes from the annual Community Survey on ICT Usage of Firms, performed by Statistics Norway. This survey includes information on the use of broadband internet in firms. In each year, the survey samples from the universe of joint-stock firms with at least five employees. The survey design is a stratified random sampling by industry and the number of employees.

The data on broadband access comes from the Norwegian Ministry of Government Administration. The ministry monitors the access of broadband internet to households, and the suppliers of broadband access to

¹²These firms cover the vast majority of revenues and workers in the private sector (about 81 % of revenues and 71 % of workers in 2001).

end-users are therefore required to file annual reports about their coverage rates to the Norwegian Telecommunications Authority. The coverage rates are based on information on the area signal range of the local access points and detailed information on the place of residence of households. In computing the coverage rates at the municipality level, it is taken into account that multiple suppliers may provide broadband access to households living in the same area, so that double counting is avoided.

Socio-economic data

Most of our socio-economic data come from administrative registers provided by Statistics Norway. Specifically, we use a rich longitudinal database which covers every resident from 1998 to 2008. It contains individual demographic information (regarding gender, age, marital status and number of children), socio-economic data (educational attainment, income, employment status), and geographic identifiers for municipality of residence. The information on educational attainment is based on annual reports from Norwegian educational establishments, whereas the income data and employment data are collected from tax records and other administrative registers. The household information is from the Central Population Register.

Hourly wages and occupation

While the employer-employee registry contains data on employment status and annual wages of all workers, it does not provide information on hourly wages (or hours of work). When looking at the impact of broadband internet on hourly wages, we use data from Statistics Norway's Wage Statistics Survey. In each year, the survey provides information on hourly wages and occupations. For employees in the private sector, the data is based on an annual stratified random sampling of all firms. The survey covers all employees in the public sector. Taken together, the information on hourly wages covers about 80 percent of Norwegian employees in every year (100 percent of the public sector employees and 70 percent of the private sector employees).

Sample selection and summary statistics

In the production function estimation, we use the sample of joint-stock firms recorded in the internet survey (which all have at least five employees). In the interest of external validity, we exclude firms that are carrying out extraction of natural resources (including oil, gas and fish).¹³ We refine this sample to be appropriate for estimation of production functions by focusing on firms with at least one employee in each of the two levels of skill. The baseline sample consists of 18,675 firms over the period 2001–2007. Throughout the paper, we use sampling weights to produce representative estimates for the corresponding population of joint-stock firms (192,312 firms).¹⁴

Table 2.1 displays summary statistics for key firm variables over time.¹⁵ The first panel displays the mean of output and non-labor inputs over time, with standard deviations in parentheses. In the production function, we use value added as the dependent variable, defined as revenues (total sales) net of intermediates (procurement of materials and intermediate inputs). We measure capital as the value of total stock of fixed assets. It is evident that these variables are fairly stable over time, perhaps with a weakly increasing trend in revenues, value added and intermediates.

¹³The production function estimates barely move if we include firms carrying out extraction of natural resources.

¹⁴The production function estimates are very similar if we instead use the unweighted sample of firms. The same is true if we exclude firms which locate in the capital (Oslo) or the three big cities.

¹⁵Throughout this paper, all monetary figures are fixed at 1998 level after adjusting for inflation. For the figures expressed in U.S. dollars (USD), we have used the following exchange rate: NOK/USD = 7.5.

Table 2.1: Descriptive statistics of baseline firm sample

	2001	2004	2007	Overall	Overall
<i>Input-output (USD, thousands)</i>	(log)	(log)	(log)	(log)	(level)
Revenues	7.45 (1.20)	7.45 (1.23)	7.72 (1.35)	7.51 (1.26)	6,190 (45,766)
Value added	6.65 (1.10)	6.68 (1.12)	6.99 (1.24)	6.74 (1.15)	2,412 (14,894)
Intermediates	6.45 (1.82)	6.42 (1.84)	6.60 (2.01)	6.47 (1.85)	3,779 (36,864)
Capital	4.54 (1.77)	4.34 (1.88)	4.50 (2.01)	4.41 (1.88)	1,112 (18,532)
<i>Wage bills (USD, thousands)</i>	(log)	(log)	(log)	(log)	(level)
Total	5.72 (1.07)	5.79 (1.10)	6.20 (1.16)	5.87 (1.12)	906 (5,261)
Unskilled	4.78 (1.28)	4.74 (1.30)	5.03 (1.35)	4.83 (1.29)	336 (2,355)
Skilled	4.92 (1.37)	5.08 (1.38)	5.57 (1.41)	5.17 (1.39)	570 (3,287)
<i>Employees</i>	(log)	(log)	(log)	(log)	(level)
Total	2.67 (0.91)	2.71 (0.91)	3.03 (0.99)	2.78 (0.93)	34 (199)
Unskilled	1.86 (1.04)	1.81 (1.05)	2.01 (1.13)	1.88 (1.06)	15 (107)
Skilled	1.88 (1.07)	2.01 (1.04)	2.40 (1.14)	2.07 (1.08)	19 (100)
<i>Number of firms</i>					
Survey	2,422	2,536	3,341		18,675
Population	26,631	27,623	27,807		192,312

Note: The baseline firm sample consists of joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. (Un)Skilled comprises workers with(out) high school diploma. Detailed descriptions of the variables are given in Appendix Table A1.

The second and third panel of Table 2.1 show means and standard deviations of wage bills and number of employees by skill levels. There is a steady increase in both the wage bills and the number of employees over time, especially for the high skilled. Following Fox and Smeets (2011), our main analysis measures labor inputs by wage bills instead of the number of workers. This has the advantage of making the measure of physical capital and human capital more comparable: Physical capital is measured in terms of monetary units to reflect the quality of the machinery employed, while using the wage bill to proxy for labor input also implies measuring labor in terms of its expense in order to better reflect its quality.¹⁶

Appendix Figure A1 displays the distribution of firms by industry. This figure shows the industry composition in our sample and in the corresponding population of firms. The four main industries are wholesale/retail, manufacturing, construction, and real estate/business services. This holds true both in terms of number of firms, aggregate value added, number of employees, and total wage bills. We can also see that the distributions in our sample (with sampling weights) closely mirror the distributions for the population of firms. The ability of our sampling weights to produce representative estimates are confirmed in Appendix Figures A2 and A3: The former displays the distributions of output and inputs across firms, while the latter

¹⁶Our findings of skill-biased technical change from broadband adoption in firms are robust to measuring labor inputs by the number of workers instead of the wage bill.

shows the time trends in these variables.

Appendix Table A2 reports estimates from a standard Cobb-Douglas production function (without broadband adoption), based on our sample (with sampling weights) and the population of firms. The first two columns report OLS estimates, while the last two columns use the method for estimating production functions proposed by Levinsohn and Petrin (2003). It is evident that whether we use our sample or the population of firms matters little for the estimated output elasticities. We can also see that our estimates align well with the findings in previous studies. As predicted by theory, OLS overstates the labor coefficients because the level of inputs chosen is positively correlated with unobserved productivity. The magnitudes of the output elasticities of capital and labor are comparable to what found in previous studies using micro data (see e.g. Pavcnik, 2002; Fox and Smeets, 2011).

Lastly, Appendix Table A3 displays summary statistics for the labor outcomes over time. When estimating the employment effects, we consider the full population of individuals between the ages of 18 and 67 (the mandatory retirement age). To estimate the impact on hourly wages, we consider all workers between the ages of 18 and 67 who are recorded in the wage statistics surveys. Due to a sluggish Norwegian economy, the employment rates decline somewhat between 2000 and 2004. By comparison, hourly wages were steadily increasing over the entire period 2001–2007.

2.4 Expansion of broadband internet

Over the past decade, many OECD countries were planning the expansion of services related to information and communications technology. In Norway, the key policy change came with the National Broadband Policy, introduced by the Norwegian Parliament in the late 1990s. This section provides details about the program and describes the expansion of broadband internet.¹⁷

The program

The National Broadband Policy had two main goals. The first was to ensure that every area of the country had access to broadband at a reasonable and uniform price. The second was to ensure that the public sector quickly adopted broadband internet.

The Norwegian government took several steps to reach these goals. First and foremost, it invested heavily in the necessary infrastructure. The investment in infrastructure was largely channeled through the (state-owned) telecom company Telenor, which was the sole supplier of broadband access to end-users in the early 2000s and continues to be the main supplier today. Moreover, virtually all broadband infrastructure was, and still is, owned and operated by Telenor.

Second, local governments were required to ensure access to broadband internet by 2005 to local public institutions, such as administrations, schools, and hospitals (St.meld.nr. 49, 2002–2003). To assist municipalities in rural areas, the federal government provided financial support through a funding program known as *Høykom*. Local governments could receive funds from this program by submitting a project plan that had to be evaluated by a program board with expert evaluations. The stated aim was to ensure broadband coverage throughout the country. Once approved, financial support was provided in the initial years of broadband access, thus making it possible for public institutions to cover relatively high initial costs.¹⁸

¹⁷Our discussion draws on Bhuller, Havnes, Leuven, and Mogstad (2013).

¹⁸During the period 1999–2005, the *Høykom* program received more than 1000 such applications and co-funded nearly 400 projects, allocating a total of 50 million Euros (NOK 400 million). From 2002, the Ministry of Education and Research co-financed another scheme (*Høykom skole*), providing financial support for broadband infrastructure in public schools. There are

Supply and demand factors

The transmission of broadband signals through fiber-optic cables required installation of local access points. Since 2000, such access points were progressively rolled out, generating considerable spatial and temporal variation in broadband coverage. The staged expansion of access points was in part due to limited public funding, but also because Norway is a large and sparsely populated country. There are often long driving distances between the populated areas, which are mostly far apart or partitioned by mountains or the fjord-gashed shoreline.¹⁹

The government reports describing the National Broadband Policy and the roll-out of broadband access points (see St.meld.nr. 38 (1997-1998); St.meld.nr. 49 (2002-2003); Bhuller, Havnes, Leuven, and Mogstad, 2013), suggest that the main *supply factors* determining the timing of roll-out are topographical features and existing infrastructure (such as roads, tunnels, and railway routes), that slow down or speed up physical broadband expansion.²⁰ Based on the program accounts, we expect the potential *demand factors* to be related to public service provision, income level, educational attainment, and the degree of urbanization in the municipality.

Descriptive statistics

Figure 2.2 displays the evolution of broadband coverage to households between 2001 and 2007. In each year, we report the overall mean broadband coverage rate and the distribution of broadband coverage rates across municipalities. There is considerable variation, both across municipalities and over time. By 2000, broadband transmission centrals were installed in the cities of Oslo, Stavanger, and Trondheim, as well as in a few neighboring municipalities of Oslo and Trondheim. However, less than one-third of the households were covered by a broadband service provider in each of these municipalities.

Appendix Figure A4 shows geographic differences in the broadband coverage rates across municipalities over time. The maps illustrate that for a large number of municipalities there was no broadband coverage in the first few years, whereas most municipalities had achieved fairly high coverage rates in 2005. Moreover, there is considerable variation in coverage rates within the municipalities in these years.

Figure 2.2 displays the evolution of broadband internet use in firms between 2001 and 2007. There is considerable variation, both across municipalities and over time. In particular, the increase in coverage to households seems to have spilled over into increased access and take-up among firms. This is mirrored by an increase in the average usage rate of firms from around 20 percent in 2001 to above 80 percent by 2007.

2.5 Empirical strategy

Randomizing broadband adoption is not feasible: We cannot in practice force firms to adopt a new technology. One can, however, think of a social experiment which randomizes broadband coverage at the municipality level. The randomization breaks the correlation between coverage rates and unobserved determinants of productivity. The intention of our IV-approach is to mimic this hypothetical experiment. Our source of

virtually no private schools in Norway.

¹⁹The Norwegian territory covers about 149,400 square miles, an area about the size of California or Germany, with around 13 % and 6 % of those regions' populations (in 2008), respectively. The country is dominated by mountainous or high terrain, as well as a rugged coastline stretching about 1,650 miles, broken by numerous fjords and thousands of islands.

²⁰The reason is that the transmission of broadband signals through fiber-optic cables required installation of local access points. In areas with challenging topography and landscapes, it was more difficult and expensive to install the local access points and the fiber-optic cables. Furthermore, the existing infrastructure mattered for the marginal costs of installing cables to extend the coverage of broadband within a municipality and to neighboring areas.

exogenous variation in broadband adoption comes from the staged installation of broadband infrastructure, generating spatial and temporal variation in broadband coverage rates.

Below, we first describe how we use the roll-out of broadband coverage to address the endogeneity of internet adoption. Next, we discuss how we follow Levinsohn and Petrin (2003) in using intermediate inputs to control for correlation between input levels and the unobserved firm-specific productivity process.

IV model

Our IV strategy exploits that firms' access to and take-up of broadband internet changed as a side-effect of the expansion of broadband internet to households. This can be represented by the following system of equations, where the second stage is

$$y_{imkt} = x'_{imkt}\beta_0 + D_{imkt}x'_{imkt}\beta_1 + \eta_m + \tau_t + \lambda_k + \varepsilon_{imkt}, \quad (2.3)$$

while the first stages are given by

$$\begin{aligned} D_{imkt} &= x'_{imkt}\delta + Z_{mt}x'_{imkt}\phi + \gamma_m + \theta_t + \sigma_k + \nu_{imkt} \\ D_{imkt}x_{1,imkt} &= x'_{imkt}\delta_1 + Z_{mt}x'_{imkt}\phi_1 + \gamma_{1,m} + \theta_{1,t} + \sigma_{1,k} + \nu_{1,imkt} \\ &\vdots = \vdots \\ D_{imkt}x_{n,imkt} &= x'_{imkt}\delta_n + Z_{mt}x'_{imkt}\phi_n + \gamma_{n,m} + \theta_{n,t} + \sigma_{n,k} + \nu_{n,imkt} \end{aligned} \quad (2.4)$$

where y is (log) value-added, D is a dummy variable for whether the firm uses broadband internet, Z is the fraction of households in a municipality with access to broadband internet, and x is a vector of n inputs (in log) and a constant term. In all cases, subscript i denotes firm, subscript k denotes industry, subscript m denotes municipality and subscript t denotes year. All equations include a full set of municipality indicators (η_m, γ_m), year indicators (τ_t, θ_t), and (4-digit) industry indicators (σ_k, λ_k). Throughout the paper, the standard errors are clustered at the municipality level and robust to heteroskedasticity.

Unobservable determinants of production that are fixed at the municipality or industry level will be controlled for through the municipality and industry indicators, just like common time shocks are absorbed by the year indicators. To check that the estimated output elasticities are not confounded by time-varying observable factors, we report results with and without a large set of time-varying controls. In particular, we show that our estimates change little if we add control variables for the potential supply and demand factors discussed above.

Assessing the IV model

Our IV model – which controls for municipality (time and industry) fixed effects – is motivated by two features of the program. First, most of the supply and demand factors tend to vary little over time. Second, the timing of the roll-out is unlikely to co-vary with key correlates of production.

To investigate whether the data is consistent with these program features, we first regress Z_{mt} on municipality, time and industry fixed effects as well as time-varying supply and demand factors. We find that 89 % of the variation in broadband coverage can be attributed to time-invariant municipality and industry characteristics and common time effects, while less than 1 % of the variation in broadband coverage can be

attributed to a large set of time-varying variables.²¹

Second, we examine the relationship between the timing of broadband roll-out and baseline municipality characteristics. To this end, we estimate the following equation

$$\Delta Z_{mt} = \eta_m + [\theta_t \times W_{m,2000}]' \psi_t + \epsilon_{mt} \quad (2.5)$$

where $\Delta Z_{mt} = Z_{mt} - Z_{m,t-1}$, θ_t is a vector of year fixed effects, and $W_{m,2000}$ includes municipality-level information from year 2000 on demography, average levels of inputs and output, industry structures and skill composition. Demography includes income level, education, share of population residing in a densely populated locality (an urbanization indicator), size of population and level of unemployment. For inputs and output we have included municipality averages of revenues, intermediates, capital stock, number of workers and wage bill. As measures of industry structure, we use number of firms, employment share in manufacturing, employment share in wholesale, and employment share in services. The last group of municipality characteristics is the skill composition, measured as the shares of wages and workers by skill level.

Appendix Figure A5 plots the estimated coefficients from the vector ψ_t for every t (and the associated 95 % confidence intervals). Our results indicate that broadband expansion is positively related to urbanization until 2002. From 2003 onwards, there appears to be no systematic relationship between the timing of the broadband expansion and this variable. But more importantly, the timing of the expansion does not correlate with baseline industry structure, the levels of output and inputs, as well as the skill composition. Nevertheless, in the robustness analysis we demonstrate that our estimates change little when allowing for different underlying time trends in firm productivity or labor outcomes across municipalities.

Extension of the IV model

The IV approach breaks the correlation between D_{imkt} and ϵ_{imkt} . However, it is still possible that the factor inputs x_{imkt} are correlated with broadband adoption and unobserved productivity, which could bias the IV estimates of equation (2.3). Following Levinsohn and Petrin (2003, hereafter LP), we take a more structural approach to address this threat to identification of the production function.²²

LP use a structural model of an optimizing firm to derive the conditions under which intermediate inputs can be used to proxy for unobserved productivity in the production function. The error term ϵ_{imkt} is assumed to be additively separable in a transmitted component (ω_{imkt}) and an i.i.d. component (χ_{imkt}). The key difference between ω_{imkt} and χ_{imkt} is that the former is a state variable, and therefore impacts the firm's decision rule, while the latter has no impact on the firm's decision. The intermediate input demand function depends on the firm-specific state variables, ω_{imkt} and k_{imkt} ,

$$a_{imkt} = g_t(\omega_{imkt}, k_{imkt}). \quad (2.6)$$

²¹The time-varying variables include demographic factors (income level, education, share of population residing in a densely populated locality, size of population, and level of unemployment), inputs and output (municipality averages of revenues, intermediates, capital stock, number of workers and wage bill), industry structure (number of firms, employment share in manufacturing, employment share in wholesale, and employment share in services) and skill composition (shares of wages and workers by skill level).

²²Levinsohn and Petrin (2003) extend on Olley and Pakes (1996) by using intermediate inputs instead of investments as a proxy for unobserved productivity. This addresses the problem that investment is zero in a non-trivial number of cases. Both studies also show how to address endogenous exit and entry of firms. Because our information on broadband adoption comes from repeated cross-sections (rather than panel data), we abstract from endogenous entry and exit.

and it must be monotonic in ω for all relevant k .²³ The monotonicity condition for intermediate inputs means that conditional on capital, profit maximizing behavior must lead more productive firms to use more intermediate inputs.

The monotonicity allows $g_t(\omega_{imkt}, k_{imkt})$ to be inverted to yield ω as a function of intermediate inputs and capital, $\omega_{imkt} = \omega_t(a_{imkt}, k_{imkt})$. By expressing the unobserved productivity variable ω_{imkt} as a function of observables, we are able to control for ω_{imkt} in the second stage equation:

$$y_{imkt} = x'_{imkt}\beta_0 + D_{imkt}x'_{imkt}\beta_1 + \eta_m + \tau_t + \lambda_k + \omega_t(a_{imkt}, k_{imkt}) + \chi_{imkt}. \quad (2.7)$$

As in Olley and Pakes (1996) and LP, we use a polynomial expansion in a and k to approximate $\omega_t(\cdot)$. By simultaneous estimation of the first stage equations in (2.4) and the second stage equation in (2.7), we obtain consistent estimates of β_{u0} , β_{s0} , β_{k1} , β_{u1} , β_{s1} , and $\Phi_t(a_{imkt}, k_{imkt}) = \beta_{k0}k_{imkt} + \omega_t(a_{imkt}, k_{imkt})$.

While these output elasticities are sufficient to assess how broadband adoption affects labor productivity, we need to identify β_{k0} to recover the full shift in production technology. Because k_{imkt} is collinear with the non-parametric function $\omega_t(a_{imkt}, k_{imkt})$, further assumptions are necessary.²⁴ Details on how β_{k0} is estimated are found in Appendix B.

2.6 Empirical results

2.6.1 Evidence on skill bias

A virtue of our research design is that it provides a transparent way of showing how the skill bias of broadband internet is identified. To this end, we begin with a graphical depiction before turning to a more detailed regression-based analysis

Graphical evidence

Figure 2.3 illustrates the variation used in the first stages by drawing a scatter plot of the broadband coverage rates of households against the broadband user rates of firms. We let the size of the circle represent the number of firms in each bin. The figure shows a strong association between the coverage and user rates. This suggests a strong impact on broadband adoption of the increase in broadband coverage from the previous year. Indeed, when regressing user rates of firms on coverage rates of household as well as municipality, industry and year fixed effects, the coefficient on the coverage rate is about 0.23 with a standard error of 0.04. It implies that a 10 percentage point increase in broadband coverage to households induces (an additional) 2.3 % of the firms to adopt broadband internet.

Figure 2.1 in the introduction provided a visual representation of the reduced form of the IV model. In the same way, Appendix Figure A6 illustrates the reduced form of the LP approach by showing how LP estimates of the period-specific output elasticities vary with the roll-out of broadband coverage. Both set of estimates are summarized in Figure 2.4, where we normalize the period-specific output elasticities so that they equal one at time -1. We can clearly see that the growth in broadband coverage is associated with an increase in the output elasticity of skilled labor, while the output elasticity of unskilled labor decreases. The

²³For simplicity, we assume as Olley and Pakes (1996) and Levinsohn and Petrin (2003) that capital is the only state variable over which the firm has control, while intermediates, labor and broadband internet are viewed as non-dynamic input factors.

²⁴ β_{k1} is identified as the interaction of capital with D_{imkt} provides independent variation. Note also that the intercept in the production function is not separately identified from the mean of $E[\omega_{imkt}|\omega_{imkt-1}]$ without some further restriction.

similarity between the estimated output elasticities across the two graphs suggest that correlation between input levels and unobserved productivity is unlikely to drive our results. It is still possible that the association between broadband coverage and labor productivity is confounded by time-varying factors. However, there is no evidence of changes in a wide range of municipality level characteristics from time -1 to time zero, including industry structure, firm characteristics, skill composition and demographic characteristics.²⁵ As a result, the estimated output-elasticities barely move when we add time-varying control variables to the regression models.

Regression results

Table 2.2 reports production function estimates based on OLS, IV and LP. In all cases, we find that broadband adoption in firms favors skilled labor by increasing its relative productivity. Indeed, the output elasticity of skilled labor increases substantially, while broadband internet seems to be a substitute for unskilled labor. By comparison, there is little if any change in the output elasticity of capital. The full set of first stage results are reported in Appendix Table A4. The first stages are strong, with large F-statistics on the excluded instruments which mean weak instrument bias is not a concern.

To interpret the magnitude of these changes in output elasticities, it is useful to compute the implied marginal productivities. When using the LP estimates to compute marginal productivities, we find that for every \$1 spent on skilled labor, the production increases by \$1.34 in the absence of broadband internet. Broadband adoption in firms increases the productivity of skilled labor by \$0.27. For unskilled labor, however, this technological change lowers the marginal productivity by \$.06; one explanation for this finding is that broadband in firms replaces tasks previously performed by unskilled labor; these workers may then be reallocated to tasks for which they have lower comparative advantage, which pushes their marginal productivity down.

The statistical tests reported in the lower panel of Table 2.2 suggest that broadband internet is strongly skill-biased.²⁶ Under standard conditions, strong skill bias implies that wages or employment rates of (un)skilled workers should decrease (increase) with broadband internet. Table 2.3 examines these implications by estimating how the expansion of broadband internet affected hourly wages and employment rates in local labor market. Specifically, we run the following wage regression which includes interactions between broadband internet and individuals' skill level:

$$y_{lmkt} = \beta_U U_{lmkt} + \beta_S S_{lmkt} + Z_{mt}(\gamma_U U_{lmkt} + \gamma_S S_{lmkt}) + X'_{lmkt} \delta + \tau_t + \eta_m + \lambda_k + \varepsilon_{lmkt} \quad (2.8)$$

where y_{lmkt} denotes (log) hourly wage or employment status of individual l in municipality m and industry k in period t , and X is a vector of controls for his or her gender and potential experience (linear and squared terms). When estimating the employment effects, we consider the full population of individuals between the ages of 18 and 67. The impact on hourly wages pertain to workers aged 18-67 who are recorded in the wage statistics surveys.

In line with the estimated changes in labor productivity, the wage regressions show that the expansion of broadband internet improves (worsens) the labor outcomes of (un)skilled workers. For instance, the estimates imply that a 10 percentage point increase in broadband coverage in a municipality raises wages of skilled workers in that local labor market by 0.14 percent. The statistical tests reported in the lower panel of Table

²⁵For brevity, we do not include these graphs, but they are available from the authors upon request.

²⁶To test for strong skill bias, we use the method of Holm (1979) to take into account the multiplicity in testing $\beta_{s1} \leq 0$ vs. $\beta_{s1} > 0$ and $\beta_{u1} \geq 0$ vs. $\beta_{u1} < 0$.

Table 2.2: Production function estimates with two levels of skill

	Output elasticities		
	OLS	IV	LP
Capital	0.108*** (0.0069)	0.113*** (0.0170)	0.232*** (0.032)
Unskilled	0.302*** (0.0131)	0.322*** (0.0248)	0.256*** (0.0214)
Skilled	0.335*** (0.0155)	0.310*** (0.0253)	0.231*** (0.0196)
Internet \times Capital	-0.0016 (0.0086)	-0.0074 (0.0248)	-0.0070 (0.0202)
Internet \times Unskilled	-0.0198 (0.0145)	-0.0553 (0.0396)	-0.0574* (0.0333)
Internet \times Skilled	0.136*** (0.0143)	0.202*** (0.0433)	0.141*** (0.0373)
Null hypothesis		P-values	
No absolute skill bias: $\beta_{s1} \leq 0$		0.000	0.000
No relative skill bias: $\frac{\beta_{s1}}{\beta_{s0}} \leq \frac{\beta_{u1}}{\beta_{u0}}$		0.000	0.001
No strong skill bias: $\beta_{s1} \leq 0$ and/or $\beta_{u1} \geq 0$		0.087	0.081

* $p < 0.10$, ** < 0.05 , *** $p < 0.01$.

Note: Baseline firm sample over the period 2001-2007, consisting of 18,675 joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. All regressions include fixed effect for year, municipality and industry. To test for strong skill bias, we use the method of Holm (1979).

2.3 support the conclusion that broadband internet is strongly skill biased.

2.6.2 Specification checks

Table 2.2 demonstrated that the evidence on skill-biased technical change is robust to whether we use intermediate inputs to control for correlation between input levels and the unobserved firm-specific productivity process. We now report results from additional specification checks to further increase the confidence in our findings.

In Table 2.4, we challenge the validity of the empirical strategy in several ways. Column 1 repeats the baseline estimates, while columns 2 and 3 include a wide range of controls for time-varying demographic and industry characteristics. When we include these covariates, we find that the estimated output elasticities are quite similar to our baseline estimates.

An important requirement for our instrument to be valid is that the expansion of broadband internet is unrelated to different underlying time trends in firm productivity across municipalities. Columns 4-6 of Table 2.4 address this threat to identification. We first interact baseline (year 2000) covariates with a linear (column 4) and quadratic time trend (column 5). By including these controls in both the first and second stages, we allow the expansion of broadband internet to be related to different underlying time trends in productivity across municipalities, depending on their initial characteristics. A limitation of these robustness checks is that baseline characteristics may not adequately capture differential time trends across municipalities. Column 6 includes linear municipality-specific time trends in both the first and second stages. It is reassuring to find that neither the IV estimates nor the LP estimates change significantly when allowing for differential time trends across municipalities.

Table 2.3: Wage regressions with two levels of skill

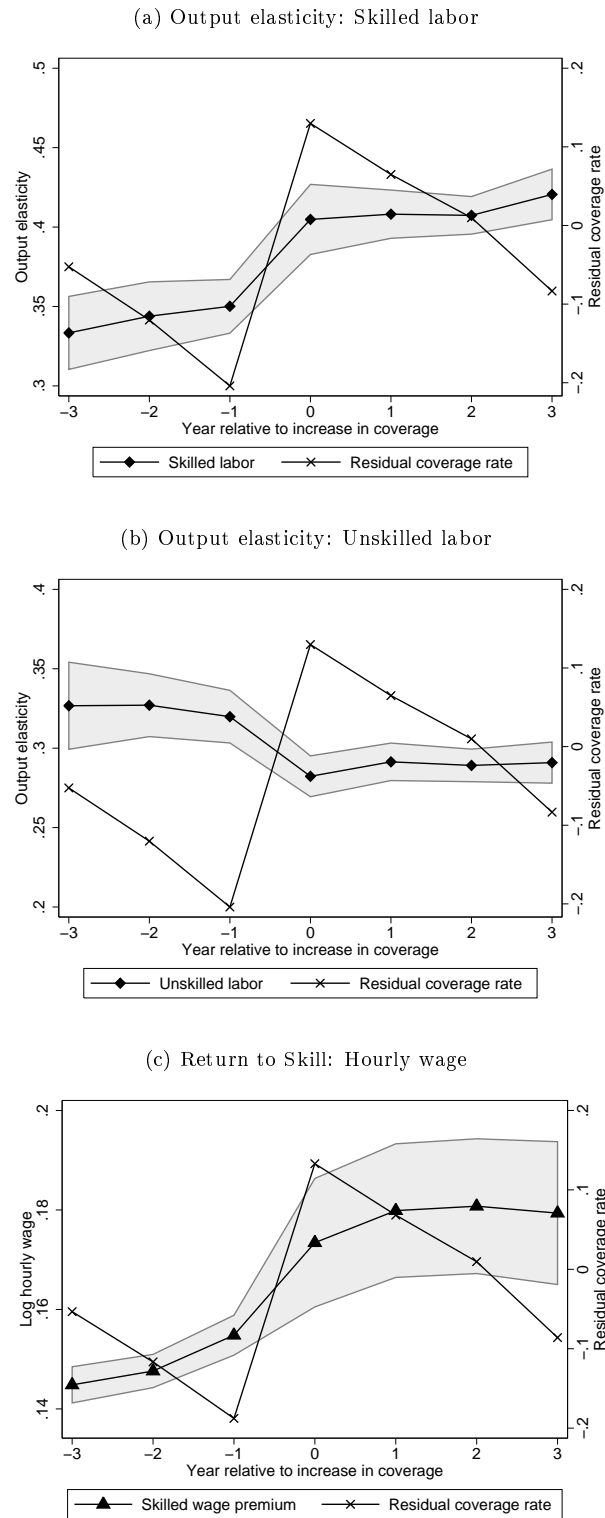
	Log hourly wage (1)	Employment rate (2)
Unskilled	9.700*** (0.0184)	0.489*** (0.0207)
Skilled	9.847*** (0.0185)	0.562*** (0.0210)
Coverage×Unskilled	-0.0281*** (0.00356)	-0.0101*** (0.00262)
Coverage×Skilled	0.0143* (0.00837)	0.0144*** (0.00491)
<hr/>		
Null hypothesis	P-values	
No absolute skill bias: $\gamma_S \leq 0$	0.044	0.002
No relative skill bias: $\frac{\gamma_S}{\beta_S} \leq \frac{\gamma_U}{\beta_U}$	0.000	0.000
No strong skill bias: $\gamma_S \leq 0$ and/or $\gamma_U \geq 0$	0.044	0.002

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: Baseline samples for labor outcomes over the period 2001-2007, consisting of all workers aged 18-67 who are recorded in the wage statistics survey (Column 1; 8,759,386 observations) and the full population of individuals between the ages of 18 and 67 (Column 2; 20,327,515 observations). Both regressions include fixed effects for year, municipality and industry and controls for gender, years of experience and years of experience squared. (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. To test for strong skill bias, we use the method of Holm (1979).


In Table 2.5, we perform the same specification checks for the wage regressions. Consistent with the baseline estimates, the alternative specifications suggest that the expansion of broadband internet improves (worsens) the labor outcomes of skilled (unskilled) workers.

Figure 2.1: Labor productivity and skill premium, pre and post expansion



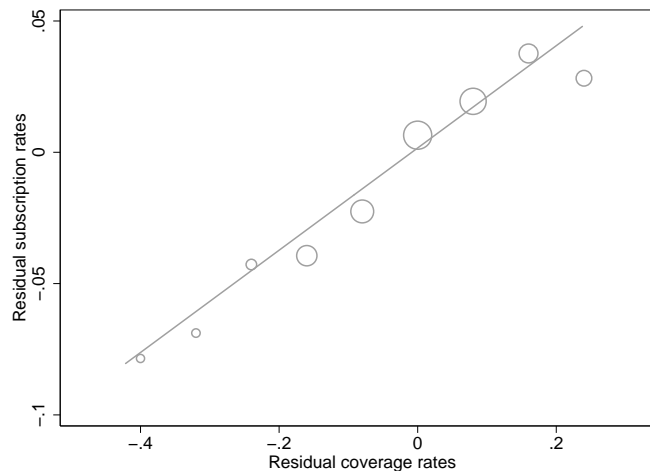
Note: Baseline samples of firms and workers during the period 2001-2007 (see Section 3 for details). Coverage rates are residuals from a regression of municipality-level broadband coverage on municipality, industry and year fixed effects. Time zero represents the year with the strongest growth in residual coverage rates in a given municipality. In each period, we estimate Cobb-Douglas production functions and wage regressions while controlling for municipality, industry and year fixed effects. Graphs (a) and (b) report period-specific OLS estimates of the output elasticity of skilled and unskilled labor. Graph (c) reports period-specific OLS estimates of log hourly wage on a dummy for skilled and controls for gender and potential experience. We define an employee as (un)skilled if he has (not) completed high school. The grey areas show 95% confidence intervals.

Figure 2.2: Broadband coverage rates of households and broadband subscription rates of firms




Note: The graphs show the overall mean and distribution of broadband coverage rates of households and broadband subscription rates of firms across municipalities for each year during the period 2001–2007.

Figure 2.3: Association between broadband subscription and coverage rates, after taking out municipality, industry and year fixed effects



Note: Baseline firm sample over the period 2001–2007, consisting of joint-stock firms recorded in the internet survey. Bins are based on the residual subscription rates of firms with bin size = 0.08. The size of the circle represents the number of firms in each bin. Sampling weights are used to ensure representative results for the population of joint-stock firms. The Y-axis reports residuals from a regression of broadband subscription rates of firms on municipality, industry and year fixed effects. The X-axis reports residuals from a regression of broadband coverage rates of households on municipality, industry and year fixed effects.

Figure 2.4: Normalized output elasticities, pre and post expansion



Note: Baseline firm sample over the period 2001-2007, consisting of joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. Coverage rates are residuals from a regression of municipality-level broadband coverage on municipality, industry and year fixed effects. Time zero represents the year with the strongest growth in residual coverage rates in a given municipality. In each time period, we estimate production functions and wage regressions while controlling for municipality, industry and year fixed effects. Graph (a) reports period-specific OLS estimates of the output elasticity of skilled and unskilled labor. Graph (b) reports period-specific LP estimates of the output elasticity of skilled and unskilled labor. In each graph, the output elasticities are normalized so that they equal one at $t = -1$.

Table 2.4: Specification checks for production function

	Time interacted with covariates						
	Baseline (1)	Covariates (2)	(3)	Linear (4)	Quadratic (5)	Linear municipality trends (6)	
<i>Panel A:</i>							
IV	Internet × Capital	-0.00741 (0.0248)	-0.0104 (0.0253)	-0.0132 (0.0248)	-0.0183 (0.0275)	-0.0187 (0.0265)	-0.00101 (0.0236)
	Internet × Unskilled	-0.0553 (0.0396)	-0.0600 (0.0406)	-0.0598 (0.0403)	-0.0677 (0.0422)	-0.0664 (0.0420)	-0.0305 (0.0431)
	Internet × Skilled	0.202*** (0.0433)	0.202*** (0.0441)	0.192*** (0.0438)	0.195*** (0.0451)	0.190*** (0.0445)	0.163*** (0.0410)
<i>Panel B:</i>							
LP	Internet × Capital	-0.00701 (0.0202)	-0.00950 (0.0205)	-0.0113 (0.0199)	-0.0188 (0.0216)	-0.0184 (0.0209)	-0.00247 (0.0202)
	Internet × Unskilled	-0.0574* (0.0333)	-0.0608* (0.0340)	-0.0624* (0.0339)	-0.0707** (0.0352)	-0.0692** (0.0348)	-0.0412 (0.0334)
	Internet × Skilled	0.141*** (0.0373)	0.140*** (0.0377)	0.132*** (0.0370)	0.134*** (0.0375)	0.129*** (0.0373)	0.105*** (0.0369)
Controls	Demographic	✓	✓	✓	✓	✓	✓
	Industry						

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: Baseline firm sample over the period 2001-2007, consisting of 18,675 joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. All regressions include fixed effect for year, municipality and industry. Column (2) adds demographic controls to the baseline model, including municipality-level information on average household income, mean years of schooling, share of population residing in a densely populated locality, size of population and level of unemployment. Column (3) also includes industry controls, consisting of municipality averages of revenues, intermediates, capital stock, number of workers and wage bills as well as employment share in manufacturing, employment share in wholesale, employment share in services, and shares of wages and workers by skill level. Columns (4) and (5) interact linear and quadratic time trends with baseline values (year 2000) of these covariates. Column (6) includes municipality-specific linear time trends.

Table 2.5: Specification checks for wage regression

	Time interacted with covariates					
	Baseline (1)	Covariates (2)	Linear (3)	Linear (4)	Quadratic (5)	Linear municipality trends (6)
<i>Panel A:</i>						
Cov × Unskilled	-0.0101*** (0.00262)	-0.00885*** (0.00221)	-0.00893*** (0.00235)	-0.0141*** (0.00330)	-0.0147*** (0.00346)	-0.0128*** (0.00338)
Cov × Skilled	0.0144*** (0.00491)	0.0155*** (0.00373)	0.0154*** (0.00359)	0.0103*** (0.00262)	0.00974*** (0.00252)	0.0116*** (0.00260)
<i>Panel B:</i>						
Cov × Unskilled	-0.0281*** (0.00356)	-0.0270*** (0.00392)	-0.0271*** (0.00410)	-0.0282*** (0.00634)	-0.0283*** (0.00636)	-0.0310*** (0.00620)
Cov × Skilled	0.0143* (0.00837)	0.0152* (0.00786)	0.0152** (0.00756)	0.0134*** (0.00498)	0.0133*** (0.00504)	0.0106** (0.00535)
Demographic		√	√	√	√	√
Industry			√	√	√	√

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: Baseline samples of labor outcomes over the period 2001-2007, consisting of all workers aged 18-67 who are recorded in the wage statistics survey (Panel B; 8,759,386 observations) and the full population of individuals between the ages of 18 and 67 (Panel A; 20,327,515 observations). All regressions include fixed effects for year, municipality and industry, and controls for gender, years of experience and years of experience squared. (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. Column (2) adds demographic controls to the baseline model, including municipality-level information on average household income, mean years of schooling, share of population residing in a densely populated locality, size of population and level of unemployment. Column (3) also includes industry controls, consisting of municipality averages of revenues, intermediates, capital stock, number of workers and wage bills as well as employment share in manufacturing, employment share in wholesale, employment share in services, and shares of wages and workers by skill level. Columns (4) and (5) interact linear and quadratic time trends with baseline values (year 2000) of these covariates. Column (6) includes municipality-specific linear time trends.

2.6.3 Understanding the pattern of adoption

Following Imbens and Angrist (1994), our IV and LP estimates should be interpreted as the local average treatment effects (LATE) for the sub-population of firms that quickly adopt broadband when it becomes accessible (often called “compliers”). To characterize these compliers and understand the pattern of broadband adoption, we estimate the following regression model separately for different types of firms:

$$D_{imkt} = \delta Z_{mt} + \gamma_m + \theta_t + \sigma_k + \nu_{imkt}. \quad (2.9)$$

We partition the baseline sample into ten mutually exclusive groups by industry and share of workers with high school diploma (above and below median within each industry). Column 1 of Table 2.6 displays the proportion of the sample in each industry–skill group. The estimates of δ for the different types of firms are shown in the second column of Table 2.6. The proportion of the compliers of a given type is then calculated as the ratio of $\hat{\delta}$ for that subgroup to the $\hat{\delta}$ in the overall sample, multiplied by the proportion of the sample in the industry–skill group. Column 3 shows the distribution of the compliers by industry and skill intensity. We see that firms with a large share of high skilled workers are overrepresented among the compliers in every industry compared to the sample of firms at large.

Columns 4–6 of Table 2.6 report the characteristics of each industry–skill group prior to the roll-out (in year 2000). Column 4 shows that in every industry the complier firms tend to be relatively large or productive (as measured by value added), while columns 5 and 6 show that they are more likely to deploy high skilled labor and use computers intensively. These findings suggest that when broadband internet becomes available, it is not randomly adopted; instead, it is more quickly adopted in firms in which complementary factors are abundant, including computers and skilled workers. This conforms to the predictions of a model of endogenous technology adoption where firms’ choices reflect principles of comparative advantage (see e.g. Beaudry and Green, 2003, 2005; Beaudry, Doms, and Lewis, 2010). The pattern of broadband adoption illustrates that the estimated effects pertain to a selective subgroup of early adopters of broadband internet; as a result, we need to be cautious in extrapolating these local estimates to the longer-run impact of broadband adoption for the population of firms at large.²⁷

²⁷For the same reason, we also need to be cautious in comparing the OLS estimates to the IV or LP estimates. The OLS estimates differ either because of endogeneity bias or because of heterogeneity across firms in the impact of adopting broadband internet.

Table 2.6: Characterizing complier firms

	Composition of sample (1)	$\hat{\delta}$ (2)	Composition of compliers (3)	Log value added (4)	Share of skilled workers (5)	Share of workers using PC (6)	Number of obs. (7)
<i>Manufacturing</i>							
Low skill	0.13	0.17	0.10	7.06	0.35	0.33	2,377
High skill	0.13	0.21	0.12	7.23	0.62	0.54	2,384
<i>Wholesale</i>							
Low skill	0.16	0.16	0.11	6.40	0.38	0.56	3,019
High skill	0.16	0.25	0.17	6.67	0.60	0.68	3,019
<i>Business services</i>							
Low skill	0.06	0.11	0.03	6.50	0.49	0.62	1,055
High skill	0.06	0.17	0.04	7.06	0.81	0.95	1,055
<i>Construction</i>							
Low skill	0.07	0.36	0.11	6.65	0.43	0.23	1,318
High skill	0.07	0.39	0.12	6.63	0.65	0.31	1,318
<i>Others</i>							
Low skill	0.08	0.18	0.07	6.50	0.38	0.27	1,565
High skill	0.08	0.21	0.08	6.50	0.62	0.49	1,565
Overall	1.00	0.23	1.00	6.68	0.50	0.50	18,675

Note: We partition the baseline firm sample into ten mutually exclusive groups by industry and share of workers with high school diploma (above and below median within each industry). Column (1) displays the proportion of the sample in each industry-skill group. Column (2) reports estimates of δ from equation (2.9) for each group. The proportion of the compliers of a given type is then calculated as the ratio of $\hat{\delta}$ for that subgroup to the $\hat{\delta}$ in the overall sample, multiplied by the proportion of the sample in the industry-skill group. Column (3) shows the distribution of the compliers by industry and skill intensity. Columns (4)-(6) report characteristics of each industry-skill group prior to the roll-out (in year 2000). Sampling weights are used to ensure representative results for the population of joint-stock firms.

2.6.4 Economic significance

To put the size of our estimates into perspective, we calculate counterfactual output elasticities and labor outcomes that would have occurred in the absence of the broadband expansion. A counterfactual outcome is measured as the actual outcome minus the predicted effect of broadband internet on the outcome (of compliers). Consider the predicted effect on the output elasticity of (un)skilled labor. In each year, we compute this as the broadband coverage rate Z_{mt} multiplied by the coefficient δ on coverage rate in equation (2.9) times the coefficient on the interaction between broadband adoption and (un)skilled labor in equation (2.7). Similarly, the predicted effect on (un)skilled wages or employment is calculated as the broadband coverage rate times the coefficient γ_S (γ_U) on the interaction between broadband coverage and (un)skilled in equation (2.8).

Panel (a) of Figure 2.5 shows the actual time trends for the estimated output elasticities, as well as the predicted counterfactual time trends. While these graphs are based on the LP estimates, the trends change little if we instead use the IV estimates. The output elasticity of skilled labor has increased over time, in part because of the expansion of broadband coverage. By comparison, unskilled workers experienced an initial decline in their labor productivity due to the expansion of broadband internet. In the last year of our sample, we predict the output elasticity of (un)skilled workers would have been 0.34 instead of 0.30 (0.23 instead of 0.24) in the absence of the broadband expansion.

Panels (b) and (c) of Figure 2.5 show similar graphs for employment rates and hourly wages. Due to a sluggish Norwegian economy, the employment rates are declining between 2000 and 2004. While the expansion of broadband internet mitigates the fall in employment rates among skilled workers, unskilled workers experience a larger decline than in the absence of the expansion. In 2005, for example, the employment rate for (un)skilled workers was 1.3 (0.9) percentage points higher (lower) than it otherwise would have been. By comparison, hourly wages increase throughout our sample period. In 2007, our estimates suggest the wages are (2.8 % lower) 1.4 % higher for (un)skilled workers than they would have been in the absence of the broadband expansion.

2.6.5 Worker heterogeneity

Our estimates of how broadband adoption affects productivity and labor outcomes are consistent with the widespread view that ICT is complementary with human capital; therefore, falls in quality-adjusted ICT prices increase skill demand and the returns to skill. We now explore what types of human capital that are complementary with broadband internet.

Since we have access to detailed information on educational attainment, we are able to make the distinction between high (post secondary education), middle (high school graduates) and low (without high school diploma) skilled labor. Moreover, we are able to differentiate between field of study among college graduates. To gain precision, we consider two types of college graduates: high return fields (consisting of science, technology, engineering, business and social sciences including law and economics) vs. low return fields (all other fields such as teaching, health, humanities, etc.). Appendix Table A5 forms the basis for the categorization of fields of study by their returns, demonstrating that workers with high vs. low return fields had systematically different wages prior to the expansion of broadband internet.

Table 2.7 reports production function estimates where we differentiate between three levels of education as well as field of study among college graduates.²⁸ Panel A suggests the skill complementarity of broadband

²⁸In Table 2.7, we refine the baseline sample to be appropriate for the estimation of these production functions by focusing on firms with at least one employee in each skill category. As a result, Panel A consists of 16,249 firms, whereas Panel B uses

Table 2.7: Production function estimates with 3 skill categories and field of study

	Production function		
	OLS	IV	LP
<i>Panel A</i>			
Internet×Unskilled	-0.0338*	-0.103**	-0.0832*
	(0.0190)	(0.0523)	(0.0427)
Internet×Medium skilled	0.0404***	0.0209	-0.0149
	(0.0137)	(0.0465)	(0.0372)
Internet×High skilled	0.0856***	0.158***	0.126***
	(0.00761)	(0.0438)	(0.0360)
<i>Panel B</i>			
Internet×Unskilled	-0.00121	-0.190**	-0.147*
	(0.0359)	(0.0929)	(0.0800)
Internet×Medium skilled	-0.00161	0.00285	-0.0259
	(0.0368)	(0.0682)	(0.0649)
Internet×High skilled low return field	0.0131	0.0735	0.0594
	(0.0163)	(0.0473)	(0.0431)
Internet×High skilled high return field	0.0797***	0.208***	0.164***
	(0.0164)	(0.0671)	(0.0620)

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: We restrict the baseline firm sample to firms with at least one employee in each skill category. As a result, Panel A (B) consists of 16,249 (11,449) joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms with at least one employee in each skill category. Unskilled comprises workers without high school diploma; medium skilled consists of high school graduates (with no post-secondary education); high skilled includes individuals with post-secondary education. High return fields comprise science, technology, engineering, and social sciences including law and economics; low return fields include teaching, health, humanities, primary sector, transport and communication (for more details, see Appendix Table A5). The standard errors are clustered at the municipality level and robust to heteroskedasticity. All regressions include fixed effect for year, municipality and industry.

internet mainly pertains to workers with post-secondary education. Panel B shows the increase in labor productivity is especially strong for college graduates in high return fields. In Table 2.8, we augment the wage regression in equation (2.8) to include three levels of skills and interactions between high skilled and field of study. Consistent with the estimated changes in output elasticities, we find the gains of broadband internet are largest for college graduates in high return fields.

2.7 Mechanisms

We explore several channels through which broadband internet can affect labor productivity and labor outcomes.

The first possible channel is that expansion of broadband internet changes demand for goods in favor of skilled labor. For example, broadband internet may facilitate e-commerce, either through online sales of final goods or online procurement of intermediate goods. To investigate this, we use information for our baseline sample of firms on whether they procure intermediate goods online or receive orders online. Using these variables as dependent variables, we estimate an IV model in which the second stage equation is

$$Y_{imkt} = \zeta D_{imkt} + \eta_m + \tau_t + \lambda_k + \varepsilon_{imkt}, \quad (2.10)$$

11,449 firms.

Table 2.8: Wage regressions with 3 skill categories and field of study

	Employment rate (1)	Log hourly wage (2)
<i>Panel A:</i>		
Cov×Unskilled	-0.00782*** (0.00251)	-0.0122*** (0.00323)
Cov×Medium skilled	0.00322 (0.00252)	-0.00921 (0.00600)
Cov×High skilled	0.0245*** (0.00787)	0.0221*** (0.00736)
<i>Panel B:</i>		
Cov×Unskilled	-0.00748*** (0.00254)	-0.0102*** (0.00373)
Cov×Medium skilled	0.00356 (0.00259)	-0.00734 (0.00666)
Cov×High skilled low return field	0.0185*** (0.00503)	0.00749** (0.00336)
Cov×High skilled high return field	0.0329*** (0.0119)	0.0478*** (0.0172)

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: Baseline samples of labor outcomes over the period 2001-2007, consisting of all workers aged 18-67 who are recorded in the wage statistics survey (Column 1; 8,759,386 observations) and the full population of individuals between the ages of 18 and 67 (Column 2; 20,327,515 observations). All regressions include fixed effects for year, municipality and industry and controls for gender, years of experience and years of experience squared. Unskilled comprises individuals without high school diploma; medium skilled consists of high school graduates (with no post-secondary education); high skilled includes individuals with post-secondary education. High return fields comprise science, technology, engineering, and social sciences including law and economics; low return fields include teaching, health, humanities, primary sector, transport and communication (for more details, see Appendix Table A5). The standard errors are clustered at the municipality level and robust to heteroskedasticity.

and the first stage equation is given by equation (2.9). Panel A of Appendix Table A6 reports the IV estimates. There is no evidence of a significant effect of broadband adoption on the probability of procuring online or receiving orders online. Another piece of evidence against this channel is provided in Appendix Table A9. In this table, we examine how broadband adoption affects labor productivity among firms in the tradeable sector where demand is given by the world market.²⁹ Consistent with the baseline results, we find that broadband adoption increases (decreases) the productivity of (un)skilled labor.

A second possible channel is that our findings are driven by changes in labor productivity and wages in firms that may have been directly affected by the expansion of broadband access. However, both the output elasticities and the wage regression results barely move if we exclude telecom firms (including Telenor) or IT consultancy companies from the estimation sample (see Appendix Tables A7 and A8). It could also be that broadband adoption coincides with general technical upgrading in firms. For example, the estimated changes in labor productivity might be due to investments in computers at the time broadband is adopted. We investigate this by estimating equations (2.10) and (2.9) with the share of workers using PC as the dependent variable in the second stage. The IV estimate reported in Panel B of Appendix Table A6 does not support

²⁹ Appendix Table A9 uses two measures of tradability. In Columns (1) and (2), we measure tradability in each 4-digit industry by dividing total levels of exports and imports by the value added of firms. In Columns (3) and (4), we follow Jensen and Kletzer (2005) in measuring tradability by the geographic concentration of an industry, defined as the Herfindahl index of employment shares across municipalities in each 4-digit industry. For both measures, we estimate the production function separately for firms in industries with values of tradability above and below the median in the baseline firm sample.

this channel.

Instead, our findings are consistent with the task approach to skill biased technological change. Starting with Autor, Levy, and Murnane (2003), a growing literature suggest that ICT substitutes for workers in performing routine tasks – more amenable to automatization – and complements workers in executing nonroutine abstract tasks. To explore this channel, we use the wage statistics surveys with information on occupation codes at the 4-digit level.³⁰ We link these occupation codes with job task requirements from the DOT data base provided by Autor and Dorn (2013). The measured job task intensities are (nonroutine) abstract tasks, routine tasks and (nonroutine) manual tasks. Appendix Table A10 provides examples of workplace activities with different task intensities.³¹

We find two pieces of evidence in support of the task based explanation of skill biased technological change. We first show that workers of different educational background indeed cluster disproportionately in occupations that require different tasks. Table 2.9 shows the occupational types, wages and task intensities for each skill category.³² We find that college graduates tend to move into professional and managerial occupations that are pervasive at abstract tasks and pay more. The sorting is especially strong among individuals with college degrees in high return fields. By comparison, low educated workers are over-represented in poorly paid occupations that involve relatively little abstract tasks, but require more routine or manual tasks.

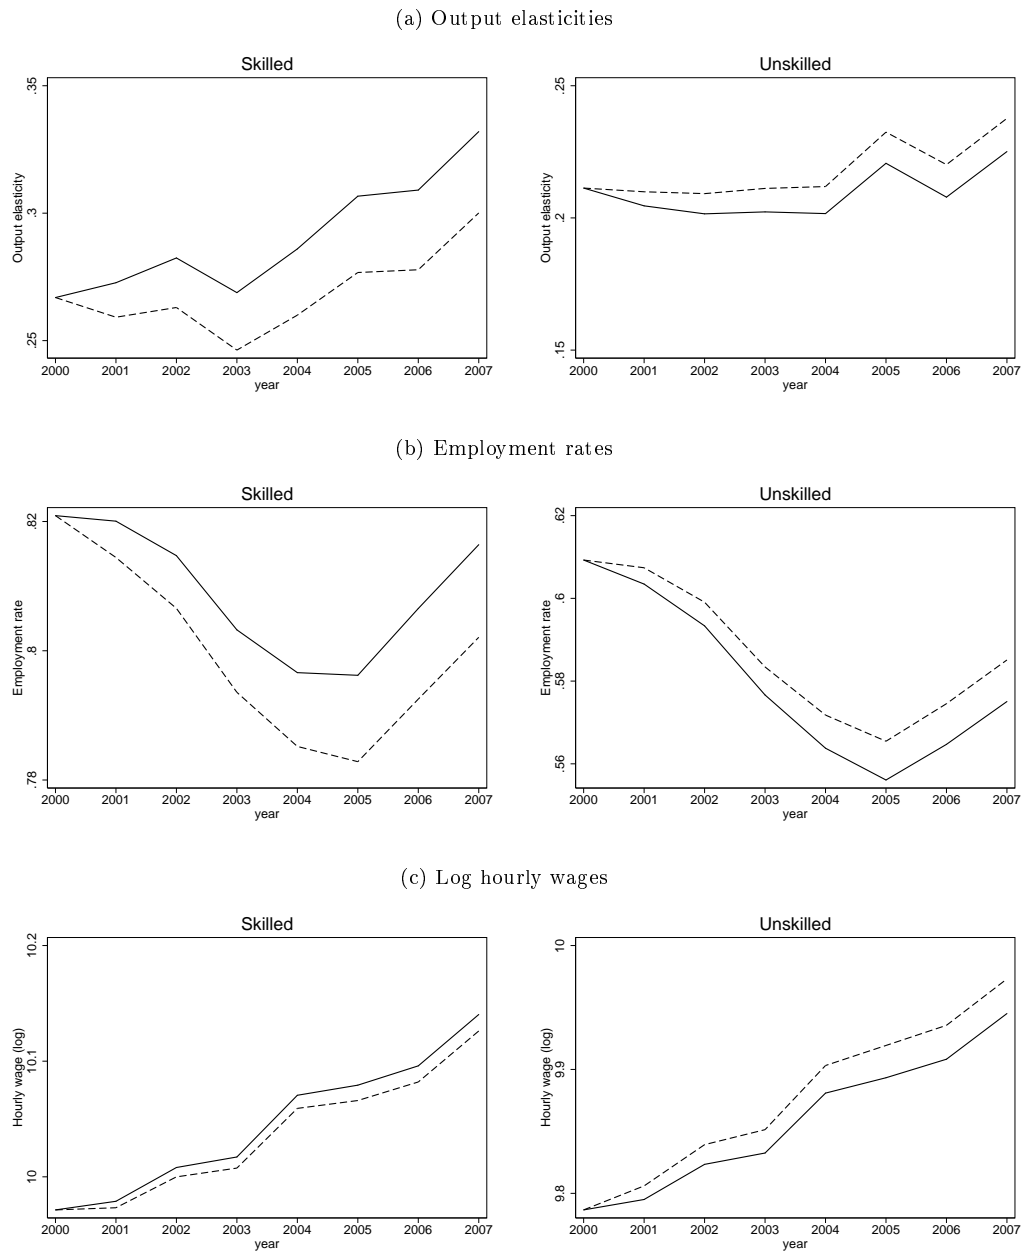
Next, we estimate wage regressions which include interactions between broadband internet and the tasks performed in jobs rather than (or in addition to) the educational credentials of workers performing those jobs. All regressions control for time, municipality, and industry fixed effects as well as (potential) experience and gender. The estimates are provided in Table 2.10. Column 1 shows the expansion of broadband internet re-enforced the wage premiums to workers performing abstract tasks. By comparison, the wages paid to jobs requiring routine tasks declined because of the broadband expansion. The estimates imply that, holding everything else equal, a 10 percentage point increase in broadband coverage in a municipality raises (lowers) hourly wages of workers with abstract (routine) task intensity at the 75th percentile by 0.9 (0.2) percent, as compared to workers at the 25th percentile of the task intensity. Columns 2-4 show the estimates are quite similar when we control for skill levels and their interaction with broadband coverage. Taken together, the results presented in Tables 2.9 and 2.10 suggest an important channel behind the skill bias of broadband internet is that it complements non-routine abstract tasks but substitutes for routine tasks whilst not affecting manual tasks.

³⁰We observe occupation codes for about 52% of the workers in the wage statistics surveys. As a result, Tables 2.9 and 2.10 are based on a subsample of the baseline sample used in the main wage regressions. When estimating the main wage regressions on this subsample, the conclusions about skill-biased technical change hold. Specifically, the estimated coefficient on the interaction between skilled and coverage rates moves from 0.0143 (s.e. = 0.008) to 0.0135 (s.e. = 0.008), whereas that for unskilled and coverage rates changes from -0.0281 (s.e. = 0.004) to -0.0423 (s.e. = 0.009).

³¹See Autor, Levy, and Murnane (2003) and Autor and Dorn (2009) for detailed descriptions of how the task intensities are measured and interpreted.

³²Following Autor and Dorn (2013), we define professional and managerial occupations as white-collar managerial, professional specialists, technical, finance, and public safety occupations.

Figure 2.5: Economic significance of changes in output elasticities and labor outcomes



Note: Solid line = actual outcome. Dashed line = counterfactual outcome in the absence of broadband internet expansion.

A counterfactual outcome is measured as the actual outcome minus the predicted effect of broadband internet on the outcome. In each year, we compute the predicted effect on the output elasticity of (un)skilled labor as the broadband coverage rate Z_{mt} multiplied by the coefficient δ on coverage rate in equation (2.9) times the coefficient on the interaction between broadband adoption and (un)skilled labor in equation (2.7). In each year, we compute the predicted effect on (un)skilled wages or employment as the broadband coverage rate times the coefficient γ_S (γ_U) on the interaction between broadband coverage and (un)skilled in equation (2.8).

Table 2.9: Occupation types, wages and task intensities by skill category.

	Occupation types			Mean relative wage			Relative task intensities		
	Professional & Managerial (1)	Other (2)	All (3)	Professional & Managerial (4)	Other (5)	All (6)	\hat{M}/\hat{A} (6)	\hat{R}/\hat{A} (7)	Number of obs.
A. 2 skill categories									
Unskilled	0.12	0.88	-0.13	0.16	-0.17	0.15	0.24	0.24	1,730,625
Skilled	0.41	0.59	0.08	0.29	-0.07	-0.09	-0.14	-0.14	2,855,708
B. 3 skill categories									
Unskilled	0.12	0.88	-0.13	0.16	-0.17	0.15	0.24	0.24	1,730,625
Medium skilled	0.17	0.83	-0.05	0.22	-0.10	0.11	0.18	0.18	1,446,022
High skilled	0.66	0.34	0.21	0.30	0.02	-0.30	-0.47	-0.47	1,409,686
C. 3 skill categories with low/high return fields									
Unskilled	0.12	0.88	-0.13	0.16	-0.17	0.15	0.24	0.24	1,730,625
Medium skilled	0.17	0.83	-0.05	0.22	-0.10	0.11	0.18	0.18	1,446,022
High skilled low return field	0.68	0.32	0.09	0.15	-0.05	0.03	-0.30	-0.30	613,708
High skilled high return field	0.65	0.35	0.30	0.43	0.07	-0.55	-0.61	-0.61	795,978

Note: We restrict the baseline worker sample to individuals for which we observe occupation code at the 4-digit level in the wage statistics surveys. The occupation codes are linked with measures of task intensity from the Dictionary of Occupational Title, as reported by Autor and Dorn (2013). Unskilled comprises individuals without high school diploma; medium skilled consists of high school graduates (with no post-secondary education); high skilled includes individuals with post-secondary education. High return fields comprise science, technology, engineering, and social sciences including law and economics; low return fields include teaching, health, humanities, primary sector, transport and communication. Columns (1) and (2) show the occupation distribution by skill category. Following Autor and Dorn (2013), we define professional and managerial occupations as white-collar managerial, professional specialists, technical, finance, and public safety occupations. Columns (3)-(5) provide mean relative wage by type of occupation and skill category. The mean relative wage is defined as the average log hourly wage within each skill-occupation group relative to the overall sample mean. Columns (6) and (7) report relative task intensities by skill level. Following Autor and Dorn (2009, 2013), we calculate \hat{M}/\hat{A} (\hat{R}/\hat{A}) as the mean of the logarithm of the ratio of the manual (routine) task content in an occupation to the abstract task content in the same occupation. Both measures are standardized with mean zero and variance one.

Table 2.10: Wage regressions with task intensities

Dep. variable: Log hourly wage	Skill categories			
	(1)	(2)	(3)	(4)
Abstract	0.364*** (0.0143)	0.326*** (0.0137)	0.265*** (0.0142)	0.271*** (0.0139)
Routine	-0.0579*** (0.00605)	-0.0627*** (0.00575)	-0.0633*** (0.00534)	-0.0583*** (0.00500)
Manual	0.0287*** (0.00785)	0.0241*** (0.00756)	0.0169** (0.00726)	0.0181*** (0.00642)
Coverage×Abstract	0.178*** (0.0325)	0.170*** (0.0293)	0.161*** (0.0300)	0.158*** (0.0291)
Coverage×Routine	-0.0347*** (0.00740)	-0.0358*** (0.00750)	-0.0340*** (0.00737)	-0.0340*** (0.00739)
Coverage×Manual	-0.00183 (0.0121)	-0.00100 (0.0117)	-0.000867 (0.0107)	0.00859 (0.00930)
Controls:				
Skill categories		√	√	√
Coverage*Skill category		√	√	√

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: We restrict the baseline worker sample to individuals for which we observe occupation code at the 4-digit level in the wage statistics surveys (4,586,333 observations). The occupation codes are linked with measures of task intensity from the Dictionary of Occupational Title (DOT), as reported by Autor and Dorn (2013). Following Autor, Levy, and Murnane (2003), we convert the DOT measures into percentiles of the task distribution. Unskilled comprises individuals without high school diploma; medium skilled consists of high school graduates (with no post-secondary education); high skilled includes individuals with post-secondary education. High return fields comprise science, technology, engineering, and social sciences including law and economics; low return fields include teaching, health, humanities, primary sector, transport and communication. Column (1) presents results from a regression of log hourly wages on task intensities and their interaction with broadband coverage in the local labor market. Column (2) adds indicator variables for two levels of skill and their interaction with broadband coverage. Column 3 includes indicator variables for three levels of skills and their interaction with broadband coverage. Column (4) includes indicator variables for three levels of skill, high return and low return field of study, and their interaction with broadband coverage. All regressions include fixed effects for year, municipality and industry, and controls for gender, years of experience and years of experience squared. Standard errors are heteroskedasticity robust and clustered at the municipality level.

2.8 Conclusion


In 2009, the U.S. Congress asked for a plan that would provide affordable broadband service to all America's citizens. In other OECD countries, there has been similar calls for public funding to roll out broadband infrastructure. While government agencies are projecting broadband penetration rates to be important for productivity and job creation, there is little scientific evidence to substantiate these claims. We contribute by examining how adoption of broadband internet in firms shifts their production technology and changes the productivity and labor outcomes of different types of workers. Our analysis employs several data sources from Norway that we can link through unique firm and individual identifiers. Using firm-level information on value added, capital, labor by skill level and broadband adoption, we estimate production functions where firms can change their technology by adopting broadband internet. We complement these production function estimates with wage regressions which include interactions between broadband internet and skill levels. A public program with limited funding rolled out broadband access points, and provides plausibly exogenous variation of broadband adoption in firms.

We find that broadband adoption favors skilled labor over unskilled labor by increasing its relative productivity. The estimated increase in productivity of skilled labor is especially large for college graduates in high return fields such as science, technology, engineering and business. By comparison, broadband internet is a substitute for workers without high school diploma, lowering their marginal productivity. Consistent with the estimated changes in labor productivity, wage regressions show that the expansion of broadband internet improves (worsens) the labor outcomes of skilled (unskilled) workers. We explore several possible explanations for the skill-bias of broadband internet. We find suggestive evidence that broadband internet complements skilled workers in executing nonroutine abstract tasks, and substitutes for unskilled workers in performing routine tasks.

Taken together, these findings have important implications for the debate about the role of government policies in encouraging productivity growth. Our estimates suggest that policy increasing the broadband penetration rates could enhance firm productivity. A related issue is why policy changes, even if they encourage productivity, do not always happen. One explanation is that established interests earning rents in the unreformed environment could be able to stave off reform. Examining who wins, who loses, and by how much could inform about the nature of these barriers and how to design compensation schemes. Our study points to the skill bias of broadband induced shift in production technology as a barrier to government investment in broadband infrastructure.

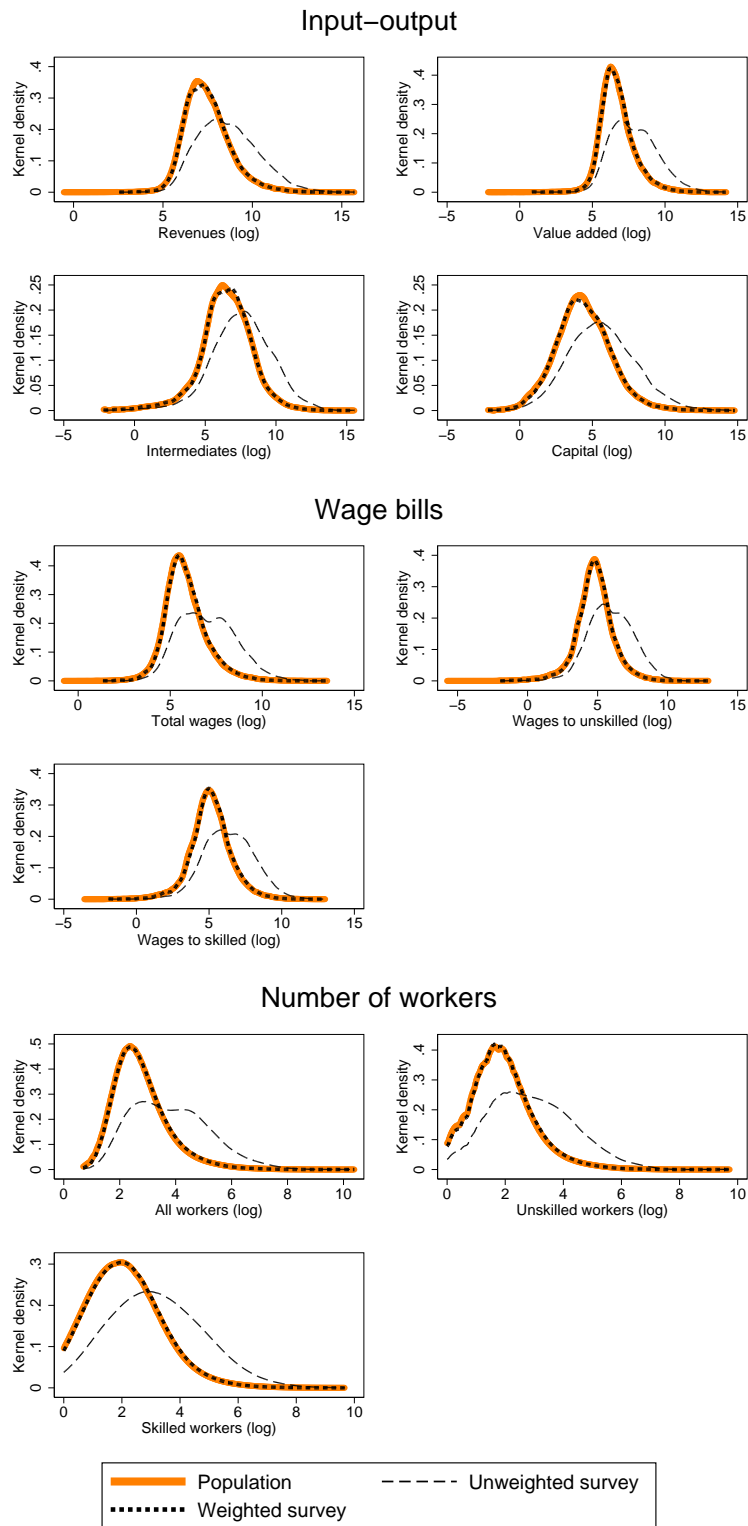
2.A Appendix A: Additional Tables and Figures

Figure A1: Distribution of firms by industry




Note: The survey sample consists of joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. We exclude firms that are carrying out extraction of natural resources (including oil, gas and fish).

Figure A2: Cross-sectional distribution of key firm variables




Note: The survey sample consists of joint-stock firms recorded in the internet survey over the period 2001-2007. The weighted survey uses sampling weights to ensure representative results for the population of joint-stock firms. We exclude firms that are carrying out extraction of natural resources (including oil, gas and fish). Detailed descriptions of the variables are given in Appendix Table A1.

Figure A3: Time trends in key firm variables



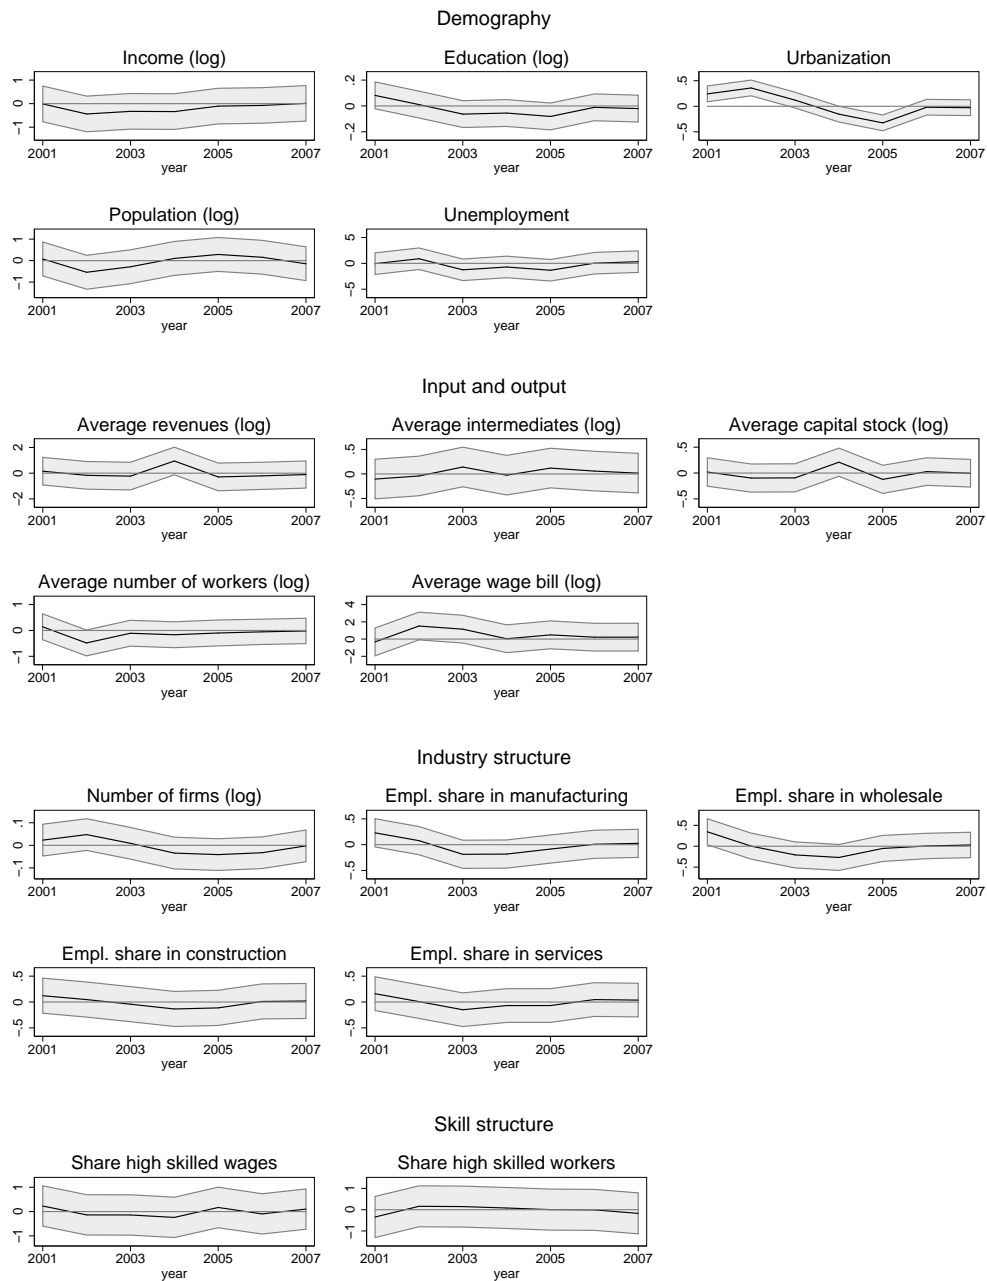
Note: The survey sample consists of joint-stock firms recorded in the internet survey over the period 2001-2007. The weighted survey uses sampling weights to ensure representative results for the population of joint-stock firms. We exclude firms that are carrying out extraction of natural resources (including oil, gas and fish). Detailed descriptions of the variables are given in Appendix Table A1

Figure A4: Geographical distribution of broadband coverage rates.



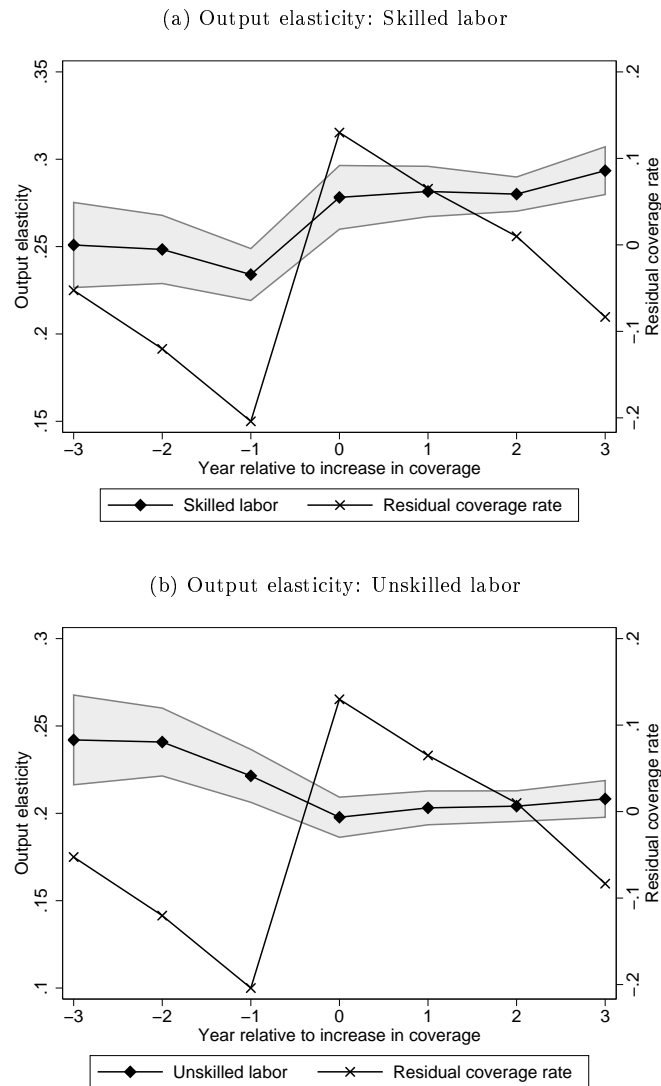
Note: The graphs show the geographical distribution of broadband coverage rates of households in 2001, 2003 and 2005.

Figure A5: Timing of broadband expansion and baseline covariates



Note: We report estimates from equation (2.5) of the vector ψ_t for every t (and the associated 95 % confidence intervals). Demography includes income level, education, share of population residing in a densely populated locality (an urbanization indicator), size of population and level of unemployment. As measures of industry structure, we use number of firms, employment share in manufacturing, employment share in wholesale, and employment share in services. For inputs and output we have included municipality averages of revenues, intermediates, capital stock, number of workers and wage bill.

Figure A6: Period-specific LP estimates of output elasticities, pre and post expansion



Note: Baseline firm sample over the period 2001-2007, consisting of joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. Coverage rates are residuals from a regression of municipality-level broadband coverage on municipality, industry and year fixed effects. Time zero represents the year with the strongest growth in residual coverage rates in a given municipality. In each time period, we estimate production functions while controlling for municipality, industry and year fixed effects. Graphs (a) and (b) report period-specific LP estimates of the output elasticity of skilled and unskilled labor. The grey areas show 95% confidence intervals.

Table A1: Variable definitions

Variable	Description
<i>Firm accounts</i>	Source: The Account Statistics.
Revenues	Total sales by a firm in year t .
Intermediates	Procurement of materials and intermediate inputs of a firm in year t .
Capital	Value of total fixed assets of a firm in year t .
Value added	Sales minus intermediates of a firm in year t .
Industry	4-digit code classifying a firm's main activity in year t according to the Standard Industry Classification (SIC2002) system.
Municipality	4-digit code for the municipality in which a firm is located in year t .
Exports	Total value of exported goods of a firm in year t .
Imports	Total value of imported goods of a firm in year t .
<i>Internet variables</i>	Source: The community survey on ICT in firms
Broadband	Dummy variable for whether a firm has adopted broadband internet (speed at or above 256 kilobits per second) in year t .
Procurement done online	Dummy variable for whether a firm conducts at least part of its procurement online in year t .
Revenues from online orders	Dummy variable for whether at least part of a firm's total revenues comes from online orders in year t .
<i>Employees</i>	Source: Register of Employers and Employees and the Wage Statistics Survey.
Annual wages	Annual pre-tax wages in year t
Employment status	Dummy variable for whether annual wages exceed the substantial gainful activity threshold in year t (USD 6,850 in 2001), which defines employment in the Social Security System.
Hourly wages	Hourly pre-tax wage per October in year t .
Occupation	4-digit occupation code of a job in year t .
<i>Individual characteristics</i>	Source: National Education Database and Central Population Register.
Education level	Years of schooling.
Field of study	Post-secondary field of study.
Municipality	Municipality of residence in year t .
Age	The age of a worker in year t .
Potential experience	Age in year t - years of schooling - 7
Gender	The gender of a worker.
<i>Internet coverage</i>	Source: Norwegian Ministry of Government Administration.
Coverage rate	Fraction of households in a given municipality with access to broadband internet (speed at or above 256 kilobits per second) in year t .
<i>Demographic controls</i>	Source: Central Population Register.
Urbanization	Population share living in densely populated area in a given municipality in year t .

Variable	Description
Income	Average annual disposable income across individuals aged 16–59 years in a given municipality in year t .
Education	Average years of schooling across individuals aged 16–59 in a given municipality in year t .
Unemployment	Unemployment rate among individuals aged 16–59 in a given municipality in year t .
<i>Industry and firm controls</i>	Source: The Account Statistics and Register of Employers and Employees.
Share of skilled workers	Share of employed workers with high school diploma in a given municipality in year t .
Share of total wages to skilled workers	Share of the total wage bill paid to workers with high school diploma in a given municipality in year t .
Share of employment by industry	Share of workers in the manufacturing/wholesale/service industry in a given municipality in year t .
Average input levels	Average level of capital stock/value added/number of workers/wages paid/revenues across firms in a given municipality in year t .

Table A2: Production function estimates without broadband internet

	OLS		LP	
	Population	Weighted Survey	Population	Weighted Survey
<i>Panel B: 1 skill category</i>				
Capital	0.0818*** (0.00297)	0.0891*** (0.00474)	0.124*** (0.020)	0.1888*** (0.018)
Labor	0.825*** (0.00299)	0.833*** (0.00741)	0.620*** (0.00418)	0.650*** (0.00991)
<i>Panel B: 2 skill categories</i>				
Capital	0.103*** (0.00333)	0.110*** (0.00555)	0.142*** (0.023)	0.225*** (0.019)
Unskilled	0.295*** (0.00508)	0.301*** (0.00898)	0.210*** (0.00375)	0.225*** (0.00703)
Skilled	0.421*** (0.00805)	0.423*** (0.0155)	0.295*** (0.00665)	0.303*** (0.0149)
Observations	192,312	18,675	192,312	18,675

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: The survey sample consists of joint-stock firms recorded in the internet survey over the period 2001-2007. Sampling weights are used to ensure representative results for the population of joint-stock firms. (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. All regressions include fixed effect for year, municipality and industry.

Table A3: Descriptive statistics of baseline sample

	2001	2004	2007	Total
<i>Employment rate</i>	(level)	(level)	(level)	(level)
Total	0.72 (0.45)	0.70 (0.46)	0.72 (0.45)	0.71 (0.45)
Unskilled	0.60 (0.49)	0.56 (0.50)	0.58 (0.49)	0.58 (0.49)
Skilled	0.82 (0.38)	0.80 (0.40)	0.82 (0.39)	0.81 (0.39)
<i>Hourly wage</i>	(log)	(log)	(log)	(log)
Total	2.95 (0.27)	3.05 (0.28)	3.12 (0.30)	3.04 (0.30)
Unskilled	2.83 (0.21)	2.92 (0.22)	2.98 (0.23)	2.91 (0.24)
Skilled	3.02 (0.28)	3.11 (0.29)	3.18 (0.30)	3.10 (0.31)
Number of individuals, aged 18-67	2,829,739	2,899,342	2,991,389	20,327,515
Number of workers in wage survey, aged 18-67	1,161,912	1,246,036	1,349,481	8,759,386

Note: (Un)Skilled comprises workers with(out) high school diploma. Detailed descriptions of the variables are given in Appendix Table A1.

Table A4: First stage regressions: 2 skill categories.

Dependent variable:		Internet	Internet × Capital	Internet × Unskilled	Internet × Skilled	
		(1)	(2)	(3)	(4)	
<i>Panel A:</i> IV	Capital	0.0237*** (0.00810)	0.223*** (0.0454)	0.294*** (0.0937)	0.283*** (0.0949)	
	Unskilled	0.0341*** (0.0122)	0.236*** (0.0585)	0.527*** (0.134)	0.485*** (0.146)	
	Skilled	0.0596*** (0.0136)	0.392*** (0.0647)	0.784*** (0.160)	0.877*** (0.168)	
	Coverage	0.589*** (0.167)	1.709** (0.806)	-0.270 (1.964)	0.393 (1.987)	
	Coverage × Capital	-0.0159* (0.00928)	0.609*** (0.0500)	-0.192* (0.109)	-0.176 (0.110)	
	Coverage × Unskilled	-0.0294** (0.0150)	-0.167** (0.0690)	0.317* (0.162)	-0.355** (0.177)	
	Coverage × Skilled	0.00536 (0.0171)	-0.114 (0.0745)	0.0206 (0.198)	0.620*** (0.202)	
	F-value (instruments)		18.3	48.8	14.3	13.0
	<i>Panel B:</i> LP	Capital	0.0589** (0.0247)	-0.227** (0.113)	0.551* (0.282)	0.504* (0.295)
Unskilled		0.0251** (0.0122)	0.169*** (0.0575)	0.409*** (0.133)	0.352** (0.146)	
Skilled		0.0493*** (0.0136)	0.310*** (0.0645)	0.643*** (0.159)	0.717*** (0.166)	
Coverage		0.596*** (0.166)	1.936** (0.790)	0.0365 (1.955)	0.824 (1.976)	
Coverage × Capital		-0.0143 (0.00910)	0.602*** (0.0479)	-0.182* (0.106)	-0.168 (0.107)	
Coverage × Unskilled		-0.0277* (0.0150)	-0.170** (0.0669)	0.329** (0.161)	-0.344* (0.176)	
Coverage × Skilled		0.00256 (0.0169)	-0.128* (0.0734)	-0.0204 (0.196)	0.571*** (0.199)	
F-value (instruments)		18.8	51.8	15.0	13.5	

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: Baseline firm sample over the period 2001-2007, consisting of 18,675 joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. Panel A (B) provides the first stage results for the IV (LP) method.

Table A5: Wage regressions with educationan levels and fields of study

Dependent variable:	Log hourly wage	
	(1)	(2)
Medium skilled	0.0886*** (0.00319)	0.0932*** (0.00385)
High skilled	0.296*** (0.00929)	0.153*** (0.0110)
High skilled interacted with:		
Humanities		0.0990*** (0.00819)
Teacher		0.0897*** (0.0134)
Social Sciences		0.223*** (0.0113)
Business		0.215*** (0.0175)
STEM		0.259*** (0.00515)
Health		0.0864*** (0.00916)
Primary		0.174*** (0.0127)
Trans/Comm.		0.0820*** (0.00805)

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: The sampel consists of workers recorded in the wage statistics survey in year 2000 (1,108,504 observations). We report results from regressions of log hourly wages in education levels and post-secondary fields of study. Unskilled comprises workers without high school diploma; medium skilled consists of high school graduates (with no post-secondary education); high skilled includes workers with post-secondary education. The reference category for education levels is unskilled workers. The reference category for field of study is “other/unspecified”. All regressions include fixed effects for municipality and industry, and controls for gender, years of experience and years of experience squared. Standard errors are heteroskedasticity robust and clustered at the municipality level.

Table A6: IV estimates of broadband adoption on E-commerce and computerization

	IV (1)	Dependent mean (2)
<i>Panel A: E-commerce:</i>		
Dep. variable: Procuring online	0.0480 (0.118)	0.50
Dep. variable: Receiving orders online	-0.00424 (0.117)	0.24
<i>Panel B: Technical upgrading</i>		
Dep. variable: Share of workers using a PC	-0.0199 (0.0850)	0.53

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: Baseline firm sample over the period 2001-2007. Sampling weights are used to ensure representative results for the population of joint-stock firms. All regressions include fixed effect for year, municipality and industry. Standard errors are heteroskedasticity robust and clustered at the municipality level. The estimated coefficient on the instrument (broadband coverage) in the first stage regression is 0.23 with a standard error of 0.04.

Table A7: Production function estimates excluding telecom firms and IT consultancy companies

		Baseline (1)	No telecom (2)	No IT consultancy (3)
IV	Internet \times Capital	-0.00741 (0.0248)	-0.00823 (0.0246)	-0.0121 (0.0243)
	Internet \times Unskilled	-0.0553 (0.0396)	-0.0547 (0.0394)	-0.0518 (0.0397)
	Internet \times Skilled	0.202*** (0.0433)	0.202*** (0.0435)	0.200*** (0.0418)
LP	Internet \times Capital	-0.00701 (0.0202)	-0.00735 (0.0199)	-0.00999 (0.0194)
	Internet \times Unskilled	-0.0574* (0.0333)	-0.0581* (0.0332)	-0.0543 (0.0335)
	Internet \times Skilled	0.141*** (0.0373)	0.140*** (0.0376)	0.137*** (0.0358)
<i>N</i>		18,675	18,520	18,192

* $p < 0.10$, ** < 0.05 , *** $p < 0.01$.

Note: Column 1 uses the baseline firm sample over the period 2001-2007, consisting of 18,675 joint-stock firms recorded in the internet survey. Column 2 excludes telecom firms (NACE code 64), whereas Column 3 excludes IT consultancy firms (NACE code 72). Sampling weights are used to ensure representative results for the corresponding population of joint-stock firms. (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. All regressions include fixed effect for year, municipality and industry.

Table A8: Wage regression estimates excluding telecom firms and IT consultancy companies

		Baseline (1)	No telecom employees (2)	No computer workers (3)
<i>Panel A:</i> Employment rate	Cov \times Unskilled	-0.0101*** (0.00262)	-0.00997*** (0.00262)	-0.00959*** (0.00261)
	Cov \times Skilled	0.0144*** (0.00491)	0.0141*** (0.00479)	0.0147*** (0.00500)
	<i>N</i>	20,327,515	20,049,183	20,247,710
<i>Panel B:</i> Log hourly wage	Cov \times Unskilled	-0.0281*** (0.00356)	-0.0286*** (0.00359)	-0.0280*** (0.00359)
	Cov \times Skilled	0.0143* (0.00837)	0.0138* (0.00817)	0.0144* (0.00838)
	<i>N</i>	8,759,386	8,627,632	8,679,582
	Experience	✓	✓	✓
	Female	✓	✓	✓

* $p < 0.10$, ** < 0.05 , *** $p < 0.01$.

Note: Column 1 uses baseline samples of labor outcomes over the period 2001-2007. Column 2 excludes workers in telecom firms (NACE code 64), whereas Column 3 excludes workers in IT consultancy firms (NACE code 72). (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. All regressions include fixed effects for year, municipality and industry and controls for gender, years of experience and years of experience squared.

Table A9: Production function estimates in tradable and non-tradable sectors

		Trade/Revenues		Geographic concentration	
		High	Low	High	Low
		(1)	(2)	(3)	(4)
<i>Panel A:</i>	Internet × Unskilled	-0.0656	-0.0219	-0.107	-0.0221
IV		(0.0596)	(0.0476)	(0.0676)	(0.0458)
	Internet × Skilled	0.323***	0.170***	0.248***	0.208***
		(0.0736)	(0.0446)	(0.0752)	(0.0401)
<i>Panel B:</i>	Internet × Unskilled	-0.0728	-0.0225	-0.0886	-0.0347
LP		(0.0521)	(0.0419)	(0.0578)	(0.0398)
	Internet × Skilled	0.257***	0.117***	0.189**	0.147***
		(0.0693)	(0.0426)	(0.0738)	(0.0364)
Mean of tradability measure		0.31	0.02	0.00025	0.00006
<i>N</i>		9,329	9,346	9,030	9,645

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: Baseline firm sample over the period 2001-2007, consisting of 18,675 joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. All regressions include fixed effect for year, municipality and industry. We use two measures of tradability. In Columns (1) and (2), we measure tradability in each 4-digit industry by dividing total levels of exports and imports by the value added of firms. In Columns (3) og (4), we follow Jensen and Kletzer (2005) in measuring tradability by the geographic concentration of an industry, defined as the Herfindahl index of employment shares across municipalities in each 4-digit industry. For both measures, we estimate the production function separately for firms in industries with values of tradability above and below the median in the baseline firm sample.

Table A10: Examples of workplace tasks

Routine	Task measures	
	Abstract	Manual
Record-keeping	Forming/testing hypotheses	Picking/sorting
Calculation	Medical diagnosis	Repetitive assembly
Repetitive customer service (e.g., bank teller)	Legal writing	Janitorial services
	Persuading/selling	Truck driving
	Managing others	

Source: Autor, Levy, and Murnane (2003).

2.B Appendix B: Estimation Details

By simultaneous estimation of equations (2.4) and (2.7), we obtain consistent estimates of β_{u0} , β_{s0} , β_{k1} , β_{u1} , β_{s1} , and $\Phi_t(a_{imkt}, k_{imkt}) = \beta_{k0}k_{imkt} + \omega_t(a_{imkt}, k_{imkt})$. While these output elasticities are sufficient to assess how broadband adoption affects labor productivity, we need to identify β_{k0} to recover the full shift in production technology. Because k_{imkt} is colinear with the non-parametric function $\omega_t(a_{imkt}, k_{imkt})$, further assumptions are necessary.

Assuming that ω_{imkt} follows a first-order Markow process, we can write

$$\omega_{imkt} = E[\omega_{imkt}|\omega_{imkt-1}] + \xi_{imkt}.$$

This simply decomposes ω_{imkt} into its conditional expectation at time $t-1$, $E[\omega_{imkt}|\omega_{imkt-1}]$, and a deviation from that expectation, ξ_{imkt} . If the capital stock is pre-determined and current investment (which will react to productivity shocks) takes one period before it comes productive, it follows that

$$E[\xi_{imkt}k_{imkt}] = 0.$$

This is the moment which LP use to identify the capital coefficient. Roughly speaking, variation in k_{imkt} conditional on ω_{imkt-1} is the exogenous variation used for identification. To operationalize this approach in a GMM context, note that given a guess on the capital coefficient β_{k0} , we can rewrite unobserved productivity as

$$\omega_{imkt}(\beta_{k0}) = \hat{\Phi}_{imkt} - \beta_{k0}k_{imkt}.$$

Given these $\omega_{imkt}(\beta_{k0})$, we compute ξ_{imkt} by non-parametrically regressing $\omega_{imkt}(\beta_{k0})$'s on $\omega_{imkt-1}(\beta_{k0})$'s and a constant term; we then form the residual

$$\xi_{imkt}(\beta_{k0}) = \omega_{imkt}(\beta_{k0}) - \hat{\Psi}(\omega_{imkt-1}(\beta_{k0}))$$

where $\hat{\Psi}(\omega_{imkt-1}(\beta_{k0}))$ are predicted values from the non-parametric regression.

The $\xi_{imkt}(\beta_{k0})$'s are used to form a sample analogue to the above moment. i.e.

$$\frac{1}{T} \frac{1}{N} \sum_t \sum_i \xi_{imkt}(\beta_{k0}) \cdot k_{imkt}.$$

We estimate β_{k0} by minimizing the GMM criterion function

$$Q(\beta_{k0}) = \min_{\beta_{k0}} \left(\frac{1}{N} \frac{1}{T_{i1}} \sum_i \sum_{t=T_{i0}}^{T_{i1}} \xi_{imkt}(\beta_{k0}) \cdot k_{imkt} \right)^2$$

where i indexing firms and T_{i0} and T_{i1} index the second and last period in which firm i is observed.

Because our baseline sample is a repeated cross-section (rather than panel data), we adjust the above estimation procedure. Exploiting the random sampling of firms, we can identify β_{k0} from the moment

$$E[\bar{\xi}_{mkt}, \bar{k}_{mkt}] = 0.$$

where the municipality average of a variable is denoted by upper bar. By applying the above procedure to

our panel data at the municipality level, we obtain the GMM criterion function

$$Q(\beta_{k0}) = \min_{\beta_{k0}} \left(\frac{1}{M} \frac{1}{T_{m1}} \sum_m \sum_{t=T_{m0}}^{T_{m1}} \sqrt{N_{mkt}} \bar{\xi}_{mkt}(\beta_{k0}) \cdot \bar{k}_{mkt} \right)^2$$

where T_{m0} and T_{m1} index the second and last period municipality m is observed and N_{mkt} is the number of firms in municipality m in period t . To obtain standard errors on β_{k0} in Table 2.2, we use bootstrap while clustering by municipality.

Chapter 3

Technology Adoption and Productivity Dispersion: Evidence from Broadband Adoption in Firms

3.1 Introduction

A large body of evidence suggest stark differences in productivity levels across firms, even within narrowly defined industries. Firms with higher total factor productivity (TFP) will produce greater amounts of output with the same set of observable inputs than lower-TFP firms. The magnitudes involved are striking and persistent. Syverson (2004b) finds that within the same industry, a firm at the 90th percentile of the TFP distribution makes almost twice as much with the same measured inputs as a firm at the 10 percentile. Regressing a firm's current TFP on its one-year lagged TFP yields autoregressive coefficients on the order of 0.6 to 0.8 (see e.g. Foster, Haltiwanger, and Syverson (2008)).

The large dispersion in TFP across firms raises several questions. Why do firms differ so much in their abilities to convert inputs into output? Can firms control the factors that influence TFP or are they purely external products of environment in which they operate? What is the role of government policies that encourage TFP growth? The key challenge to answer these questions is that TFP is, at its heart, an unobservable residual: it is the variation in output that cannot be explained based on observable inputs like standard labor or capital measures.

A growing empirical literature aims at putting a face on that residual.¹ The existing evidence points to the importance of factors that directly impact productivity by operating within the firm. One such factor is the adoption of new technology, which could induce productivity changes and potentially explain dispersion in TFP across firms. However, estimating how adoption of new technology affects productivity has proven difficult for several reasons (see e.g. Draca, Sadun, and Van Reenen (2007)). It is often difficult to access data on technology adoption which can be linked with firm-level information on value added and factor inputs. Another key challenge is the likelihood that some determinants of production are unobserved to the econometrician but observed by the firm; if adoption of new technology depends on these determinants, then OLS estimates of production functions will be biased. On top of this, the use of inputs such as capital and

¹For a recent review, see Syverson (2011).

labor could also be correlated with technology adoption and unobserved productivity, and therefore create bias in OLS estimates.

In this paper, we estimate how the adoption of new technology induces productivity changes and explains dispersion in TFP across firms. Our context is the adoption of broadband internet in Norwegian firms over the period 2001-2007.² Many have hypothesized that information and communication technology (ICT), like broadband internet, is important in explaining productivity patterns over time and across firms. However, the evidence base is scarce and existing research has largely focused on demonstrating positive associations of ICT with productivity. Draca, Sadun, and Van Reenen (2007), in their Handbook of Information and Communication Technologies chapter, review the literature carefully and conclude that “none of the literature has produced convincing evidence of a causal impact of ICT on productivity, for example by analyzing a natural experiment”.

Our analysis employs rich data at the firm level with information on value added, capital, labor by skill level, and broadband adoption. We begin by estimating a standard production function where any output driven by broadband internet will show up as TFP. Next, we estimate a production function where firms can change their technology by adopting broadband internet. We allow for the possibility that broadband adoption is factor-neutral or factor-biased, and therefore, directly or indirectly affect output. As suggested by Syverson (2011), we use the decline in TFP dispersion from this extension of the production function as a metric of broadband internet’s importance in explaining productivity.

As a source of exogenous variation in broadband adoption, we follow Akerman, Gaarder, and Mogstad (2013) in exploiting a public program aimed at ensuring broadband access at a reasonable price to all households throughout the country.³ Because of limited funding, access to broadband was progressively rolled out, so that the necessary infrastructure (access points) was established in different municipalities at different times. We use this spatial and temporal variation in the availability of broadband across municipalities as an instrumental variable for adoption of broadband in firms. Our identification strategy – which controls for municipality (time and industry) fixed effects – is motivated by two features of the broadband program. First, most of the supply and demand factors tend to vary little over time. Second, the timing of the roll-out is unlikely to co-vary with the key correlates of productivity and labor outcomes. We demonstrate that the data is consistent with these program features, and further challenge the validity of the instrument by probing the stability of the estimates to alternative specifications, finding little cause for worry. Following Levinsohn and Petrin (2003), we also make sure that the results are robust to using intermediate inputs to proxy for unobserved productivity in the production function.

We find that broadband internet accounts for about 2 to 3 percent of the standard deviation in TFP across firms. Our results suggest the upper part of the TFP distribution is most influenced by taking broadband internet into account in the estimation of the production function. Following Olley and Pakes (1996), we decompose aggregate TFP in the economy into an unweighted average of firm-level TFP and the covariance between market share and firm-level TFP. We find that the contribution of the covariance term to aggregate TFP is reduced substantially when we account for the variation in output that can be explained by broadband adoption in firms. This suggests that broadband internet accounts for some of the allocation of production to more productive firms.

²Throughout the paper, we focus on broadband internet, defined as internet connections with download speeds that exceed 256 kbit/s. Before the expansion of broadband internet, all firms with a telephone connection would have dial-up access to internet, but limited to a bitrate of less than 56 kbit/s. Broadband internet facilitated internet use without excessive waiting times.

³See also Bhuller, Havnes, Leuven, and Mogstad (2013) who use the roll-out of broadband internet in Norway to study how internet use affects sex crimes.

Taken together, our findings have important implications for the debate about what explains productivity differences. Our results are consistent with previous studies in pointing to the importance of factors that are internal to the firm. However, this does not imply that there is limited scope for policy to influence productivity. On the contrary, our estimates suggest that policy increasing the availability of broadband internet induces differences in firms' choice of technology, which can be important in explaining the dispersion of TFP. This illustrates that even if firms control the key factors of productivity, the external products of the operating environment are important in influencing the choices firms make.

Our paper contributes to a growing empirical literature on the determinants of productivity differences across firms. Research has linked productivity levels to a number of external factors of demand and market structure. Examples include the effect of competition (Syverson (2004a); Schmitz Jr (2005)), the impact of deregulation or liberalized trade (Olley and Pakes (1996); Pavcnik (2002); Bloom, Draca, and Van Reenen (2011)), and the interaction between product market rivalry and technology spillovers (Bloom, Schankerman, and Van Reenen (2013)). Another line of work has looked at the link between firm productivity and internal factors like workforce and organizational structure. This research considers the importance of workers' human capital (Abowd, Haltiwanger, Jarmin, Lane, Lengermann, McCue, McKinney, and Sandusky (2005); Fox and Smeets (2011)), pay structure and human resources practices (Lazear (2000); Ichniowski and Shaw (2003)), managerial practices (Bloom and Van Reenen (2007)), and organizational form (Hortaçsu and Syverson (2007); Garicano and Heaton (2010)). Our paper complements this literature by providing novel evidence on how the adoption of new technology induces productivity changes and explains dispersion in TFP across firms.

Our paper is also related to a literature on the skill-bias of information and communication technology (ICT). Evidence suggest that ICT is complementary with human capital.⁴ Technological changes that reduce quality-adjusted ICT prices – such as the arrival of broadband internet – should therefore increase skill demand and returns to skill. While the labor market consequences of investments in computers and R&D have been studied extensively,⁵ there is little evidence on the skill bias of broadband adoption. A notable exception is the study by Akerman, Gaarder, and Mogstad (2013). They examine the impact of the expansion of broadband internet in Norway on the labor productivity and labor outcomes of different types of workers. Their findings suggest that broadband adoption favors skilled labor by increasing its relative productivity. Our study uses the same source of exogenous variation in broadband adoption and estimates the same type of production functions. Unlike Akerman, Gaarder, and Mogstad (2013), we explore the impact of broadband internet on productivity differences across firms.

The paper unfolds as follows. Section 2 presents a framework for estimation of productivity differences with technology choice and discusses the threats to identification in the context of previous research. Section 3 describes our data, before Section 4 discusses the expansion of broadband internet. Section 5 presents our empirical strategy, and Section 6 describes our empirical findings. The final section offers some concluding remarks.

⁴Acemoglu (2003), Bond and Van Reenen (2007) and Goldin and Katz (2007) review the extensive literature on technology-skill complementarity.

⁵See e.g. Krueger (1993), Berman, Bound, and Griliches (1994), DiNardo and Pischke (1997), Autor, Katz, and Krueger (1998), Machin and Van Reenen (1998), Beaudry and Green (Beaudry and Green (2003); Beaudry and Green (2005)) and Beaudry, Doms, and Lewis (2010).

3.2 Productivity differences with technology choice

3.2.1 A framework for estimation

Economic theory views the production technology as a function describing how a collection of factor inputs can be transformed into output, and it defines TFP as the variation in output that cannot be explained based on observable inputs. Our goal is to estimate how the adoption of new technology induces productivity changes and explains dispersion in TFP across firms.

Consider a standard Cobb Douglas production function:

$$Y_{it} = e^{\alpha} K_{it}^{\beta_k} U_{it}^{\beta_u} S_{it}^{\beta_s} e^{\nu_{it}}, \quad (3.1)$$

where Y_{it} represents output of firm i in period t , K_{it} , U_{it} and S_{it} are inputs of capital, unskilled, and skilled labor, respectively, e^{α} is the average TFP, and $e^{\nu_{it}}$ is the time- and firm-specific deviation from average TFP. It is clear from equation (3.1) that any output driven by different choice of technology, like broadband internet, will show up as TFP.

Consider instead a Cobb-Douglas production technology with exponents that potentially change with the adoption of broadband internet:

$$Y_{it} = e^{\alpha_0 + D_{it}\alpha_1} K_{it}^{\beta_{k0} + D_{it}\beta_{k1}} U_{it}^{\beta_{u0} + D_{it}\beta_{u1}} S_{it}^{\beta_{s0} + D_{it}\beta_{s1}} e^{\varepsilon_{it}}. \quad (3.2)$$

Equation (3.2) extends on the usual Cobb-Douglas technology in equation (3.1) through the indicator variable D_{it} , which is equal to one if firm i has broadband internet in period t (and zero otherwise). This simple extension allows for the possibility that broadband adoption is factor-neutral ($\alpha_1 \neq 0$) and factor-biased ($\beta_{k1} \neq 0$, $\beta_{u1} \neq 0$ or $\beta_{s1} \neq 0$), and therefore, directly or indirectly affect output.

As suggested by Syverson (2011), we use the decline in TFP dispersion from this extension of the production function as a metric of broadband internet's importance in explaining productivity. This amounts to comparing the dispersion in $e^{\nu_{it}}$ and $e^{\varepsilon_{it}}$ across firms, telling us the extent to which broadband adoption accounts for the variation in output that cannot be explained by differences in capital and labor inputs.

3.2.2 Empirical challenges and previous research

To estimate how broadband adoption shifts the production technology, it is useful to transform (3.1) and (3.2) to linear production functions:

$$y_{it} = \alpha + \beta_k k_{it} + \beta_u u_{it} + \beta_s s_{it} + v_{it}, \quad (3.3)$$

and

$$y_{it} = \alpha_0 + \beta_{k0} k_{it} + \beta_{u0} u_{it} + \beta_{s0} s_{it} + D_{it} [\alpha_1 + \beta_{k1} k_{it} + \beta_{s1} s_{it} + \beta_{u1} u_{it}] + \varepsilon_{it}, \quad (3.4)$$

where lower-case letters refer to natural logarithms. While there exists a large literature on the identification and estimation of production functions,⁶ there is little direct evidence on how adoption of new technology changes the productivity of firms.

⁶For recent reviews, see Draca, Sadun, and Van Reenen (2007) and Van Beveren (2012).

One reason is the difficulty in accessing data on technology adoption which can be linked with firm-level information on value added, capital, and labor by skill level. For instance, typical annual accounts data reports labor costs as average (or total) wages for all workers, rather than wages by the skill level of the workers (Van Beveren (2012)). Moreover, statistical agencies have only recently started to systematically collect ICT information at the firm level (Draca, Sadun, and Van Reenen (2007)). In most cases, this information relates to broad measures of expenditure or usage of ICT, rather than precise measures of specific technological changes.

Another reason is the threats to identification of production functions. A key concern is the likelihood that some determinants of production are unobserved to the econometrician but observed by the firm; if adoption of new technology depends on unobserved productivity, then OLS estimates of (3.4) will be biased. Several studies illustrate the difficulty in drawing credible inferences absent an appropriate instrumental variable. DiNardo and Pischke (1997) suggest that computer users possess unobserved skills which might have little to do with computers but which raise their productivity. Beaudry and Green (Beaudry and Green (2003); Beaudry and Green (2005)) and Beaudry, Doms, and Lewis (2010) show the speed and extent of computer adoption are likely endogenous and reflect principles of comparative advantage. Another generic concern in estimation of production functions is that the levels of capital and labor chosen could be correlated with unobserved factors of productivity. Olley and Pakes (1996) and Levinsohn and Petrin (2003) develop methods for estimating production functions that control for correlations between the use of factor inputs and unobserved productivity. However, most studies of technological change in production ignore these threats to identification and simply estimate a production function using OLS (Draca, Sadun, and Van Reenen (2007)). In particular, there have been few attempts to find instruments that affect the decision to adopt new technology.

Because we have data with firm-level information on value added, capital, labor by skill level and broadband adoption, we are able to estimate production functions like (3.4); as a source of exogenous variation in broadband adoption, we exploit that the necessary infrastructure was established in different areas at different times; and following Levinsohn and Petrin (2003), we use intermediate inputs to proxy for unobserved productivity in the production function. This enables us to address the threats to identification and provide novel evidence on how the adoption of broadband internet induces productivity changes and explains dispersion in TFP across firms.

3.3 Data

Our analysis uses several data sources, which we can link through unique identifiers for each firm, employee, and municipality. The coverage and reliability of Norwegian data are rated as exceptional in international quality assessments (see e.g. Atkinson, Rainwater, and Smeeding (1995)). Below we describe our data and sample selection, while details about the data sources and each of the variables are given in Appendix Table A1.

Firm and worker data

Our firm data comes from administrative registers, which are updated annually by Statistics Norway and verified by the Norwegian Tax Authority. The data comprises all non-financial joint-stock firms over the period

1999–2008.⁷ It contains detailed information from the firm’s balance sheets on output (such as revenues) and inputs (such as capital, labor, intermediates) as well as 4-digit industry codes and geographical identifiers at the municipality level.

We merge the firm data set with a linked employer–employee registry that contains complete records of all firms and workers for the period 1999–2008. For every employee, we know his or her length and field of education, and annual labor income. We define an employee as skilled if he or she has completed high school, while employees with less schooling are defined as unskilled.

Internet data

For the period 2001–2007, we have (i) data on broadband subscription for a stratified random sample of firms, and (ii) complete records of the fraction of households that have access to broadband internet in every municipality. As explained in detail below, we will use the former to measure broadband adoption in firms, while the latter will form the basis for our instrumental variable. Throughout this paper, broadband is defined as internet connections with download speed that exceeds 256 kbit/s.

Our data on broadband subscriptions of firms comes from the annual Community Survey on ICT Usage of Firms, performed by Statistics Norway. This survey includes information on the use of broadband internet in firms. In each year, the survey samples from the universe of joint-stock firms with at least five employees. The survey design is a stratified random sampling by industry and the number of employees.

The data on broadband access comes from the Norwegian Ministry of Government Administration. The ministry monitors the access of broadband internet to households, and the suppliers of broadband access to end-users are therefore required to file annual reports about their coverage rates to the Norwegian Telecommunications Authority. The coverage rates are based on information on the area signal range of the local access points and detailed information on the place of residence of households. In computing the coverage rates at the municipality level, it is taken into account that multiple suppliers may provide broadband access to households living in the same area, so that double counting is avoided.

Socio-economic data

Most of our socio-economic data come from administrative registers provided by Statistics Norway. Specifically, we use a rich longitudinal database which covers every resident from 1998 to 2008. It contains individual demographic information (regarding gender, age, marital status and number of children), socio-economic data (educational attainment, income, employment status), and geographic identifiers for municipality of residence. The information on educational attainment is based on annual reports from Norwegian educational establishments, whereas the income data and employment data are collected from tax records and other administrative registers. The household information is from the Central Population Register.

Sample selection and summary statistics

In the production function estimation, we use the sample of joint-stock firms recorded in the internet survey (which all have at least five employees). In the interest of external validity, we exclude firms that are carrying out extraction of natural resources (including oil, gas and fish).⁸ We refine this sample to be appropriate for

⁷These firms cover the vast majority of revenues and workers in the private sector (about 81 % of revenues and 71 % of workers in 2001).

⁸The production function estimates barely move if we include firms carrying out extraction of natural resources.

Table A1: Descriptive statistics of baseline firm sample

	2001	2004	2007	Overall	Overall
<i>Input-output (USD, thousands)</i>	(log)	(log)	(log)	(log)	(level)
Revenues	7.45 (1.20)	7.45 (1.23)	7.72 (1.35)	7.51 (1.26)	6,190 (45,766)
Value added	6.65 (1.10)	6.68 (1.12)	6.99 (1.24)	6.74 (1.15)	2,412 (14,894)
Intermediates	6.45 (1.82)	6.42 (1.84)	6.60 (2.01)	6.47 (1.85)	3,779 (36,864)
Capital	4.54 (1.77)	4.34 (1.88)	4.50 (2.01)	4.41 (1.88)	1,112 (18,532)
<i>Wage bills (USD, thousands)</i>	(log)	(log)	(log)	(log)	(level)
Total	5.72 (1.07)	5.79 (1.10)	6.20 (1.16)	5.87 (1.12)	906 (5,261)
Unskilled	4.78 (1.28)	4.74 (1.30)	5.03 (1.35)	4.83 (1.29)	336 (2,355)
Skilled	4.92 (1.37)	5.08 (1.38)	5.57 (1.41)	5.17 (1.39)	570 (3,287)
<i>Employees</i>	(log)	(log)	(log)	(log)	(level)
Total	2.67 (0.91)	2.71 (0.91)	3.03 (0.99)	2.78 (0.93)	34 (199)
Unskilled	1.86 (1.04)	1.81 (1.05)	2.01 (1.13)	1.88 (1.06)	15 (107)
Skilled	1.88 (1.07)	2.01 (1.04)	2.40 (1.14)	2.07 (1.08)	19 (100)
<i>Number of firms</i>					
Survey	2,422	2,536	3,341		18,675
Population	26,631	27,623	27,807		192,312

Note: The baseline firm sample consists of joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. (Un)Skilled comprises workers with(out) high school diploma. Detailed descriptions of the variables are given in Appendix Table A1.

estimation of production functions by focusing on firms with at least one employee in each of the two levels of skill. The baseline sample consists of 18,675 firms over the period 2001–2007. Throughout the paper, we use sampling weights to produce representative estimates for the corresponding population of joint-stock firms (192,312 firms).⁹

Table A1 displays summary statistics for key firm variables over time.¹⁰ The first panel displays the mean of output and non-labor inputs over time, with standard deviations in parentheses. In the production function, we use value added as the dependent variable, defined as revenues (total sales) net of intermediates (procurement of materials and intermediate inputs). We measure capital as the value of total stock of fixed assets. It is evident that these variables are fairly stable over time, perhaps with a weakly increasing trend in revenues, value added and intermediates.

The second and third panel of Table A1 show means and standard deviations of wage bills and number of employees by skill levels. There is a steady increase in both the wage bills and the number of employees over time, especially for the high skilled. Following Fox and Smeets (2011), we measure labor inputs by wage bills

⁹The production function estimates are very similar if we instead use the unweighted sample of firms. The same is true if we exclude firms which locate in the capital (Oslo) or the three big cities.

¹⁰Throughout this paper, all monetary figures are fixed at 1998 level after adjusting for inflation. For the figures expressed in U.S. dollars (USD), we have used the following exchange rate: NOK/USD = 7.5.

instead of number of workers. This has the advantage of making the measure of physical capital and human capital more comparable: Physical capital is measured in terms of monetary units to reflect the quality of the machinery employed, while using the wage bill to proxy for labor input also implies measuring labor in terms of its expense in order to better reflect its quality.¹¹

Appendix Figure A1 displays the distribution of firms by industry. This figure shows the industry composition in our sample and in the corresponding population of firms. The four main industries are wholesale/retail, manufacturing, construction, and real estate/business services. This holds true both in terms of number of firms, aggregate value added, number of employees, and total wage bills. We can also see that the distributions in our sample (with sampling weights) closely mirror the distributions for the population of firms. The ability of our sampling weights to produce representative estimates are confirmed in Appendix Figures A2 and A3: The former displays the distributions of output and inputs across firms, while the latter shows the time trends in these variables.

3.4 Expansion of broadband internet

Over the past decade, many OECD countries were planning the expansion of services related to information and communications technology. In Norway, the key policy change came with the National Broadband Policy, introduced by the Norwegian Parliament in the late 1990s. This section provides details about the program and describes the expansion of broadband internet.¹²

The program

The National Broadband Policy had two main goals. The first was to ensure that every area of the country had access to broadband at a reasonable and uniform price. The second was to ensure that the public sector quickly adopted broadband internet.

The Norwegian government took several steps to reach these goals. First and foremost, it invested heavily in the necessary infrastructure. The investment in infrastructure was largely channeled through the (state-owned) telecom company Telenor, which was the sole supplier of broadband access to end-users in the early 2000s and continues to be the main supplier today. Moreover, virtually all broadband infrastructure was, and still is, owned and operated by Telenor.

Second, local governments were required to ensure access to broadband internet by 2005 to local public institutions, such as administrations, schools, and hospitals (St.meld.nr. 49, 2002–2003). To assist municipalities in rural areas, the federal government provided financial support through a funding program known as *Høykom*. Local governments could receive funds from this program by submitting a project plan that had to be evaluated by a program board with expert evaluations. The stated aim was to ensure broadband coverage throughout the country. Once approved, financial support was provided in the initial years of broadband access, thus making it possible for public institutions to cover relatively high initial costs.¹³

¹¹Our findings of skill-biased technical change from broadband adoption in firms are robust to measuring labor inputs by the number of workers instead of the wage bill.

¹²Our discussion draws on Bhuller, Havnes, Leuven, and Mogstad (2013) and Akerman, Gaarder and Mogstad (2013).

¹³During the period 1999–2005, the *Høykom* program received more than 1000 such applications and co-funded nearly 400 projects, allocating a total of 50 million Euros (NOK 400 million). From 2002, the Ministry of Education and Research co-financed another scheme (*Høykom skole*), providing financial support for broadband infrastructure in public schools. There are virtually no private schools in Norway.

Supply and demand factors

The transmission of broadband signals through fiber-optic cables required installation of local access points. Since 2000, such access points were progressively rolled out, generating considerable spatial and temporal variation in broadband coverage. The staged expansion of access points was in part due to limited public funding, but also because Norway is a large and sparsely populated country. There are often long driving distances between the populated areas, which are mostly far apart or partitioned by mountains or the fjord-gashed shoreline.¹⁴

The government reports describing the National Broadband Policy and the roll-out of broadband access points (see St.meld.nr. 38 (1997-1998); St.meld.nr. 49 (2002-2003); Bhuller, Havnes, Leuven, and Mogstad (2013)), suggest that the main *supply factors* determining the timing of roll-out are topographical features and existing infrastructure (such as roads, tunnels, and railway routes), that slow down or speed up physical broadband expansion.¹⁵ Based on the program accounts, we expect the potential *demand factors* to be related to public service provision, income level, educational attainment, and the degree of urbanization in the municipality.

Descriptive statistics

Figure A1 displays the evolution of broadband coverage to households between 2001 and 2007. In each year, we report the overall mean broadband coverage rate and the distribution of broadband coverage rates across municipalities. There is considerable variation, both across municipalities and over time. By 2000, broadband transmission centrals were installed in the cities of Oslo, Stavanger, and Trondheim, as well as in a few neighboring municipalities of Oslo and Trondheim. However, less than one-third of the households were covered by a broadband service provider in each of these municipalities.

Appendix Figure A4 shows geographic differences in the broadband coverage rates across municipalities over time. The maps illustrate that for a large number of municipalities there was no broadband coverage in the first few years, whereas most municipalities had achieved fairly high coverage rates in 2005. Moreover, there is considerable variation in coverage rates within the municipalities in these years.

Figure A1 displays the evolution of broadband internet use in firms between 2001 and 2007. There is considerable variation, both across municipalities and over time. In particular, the increase in coverage to households seems to have spilled over into increased access and take-up among firms. This is mirrored by an increase in the average usage rate of firms from around 20 percent in 2001 to above 80 percent by 2007.

3.5 Technological change in production

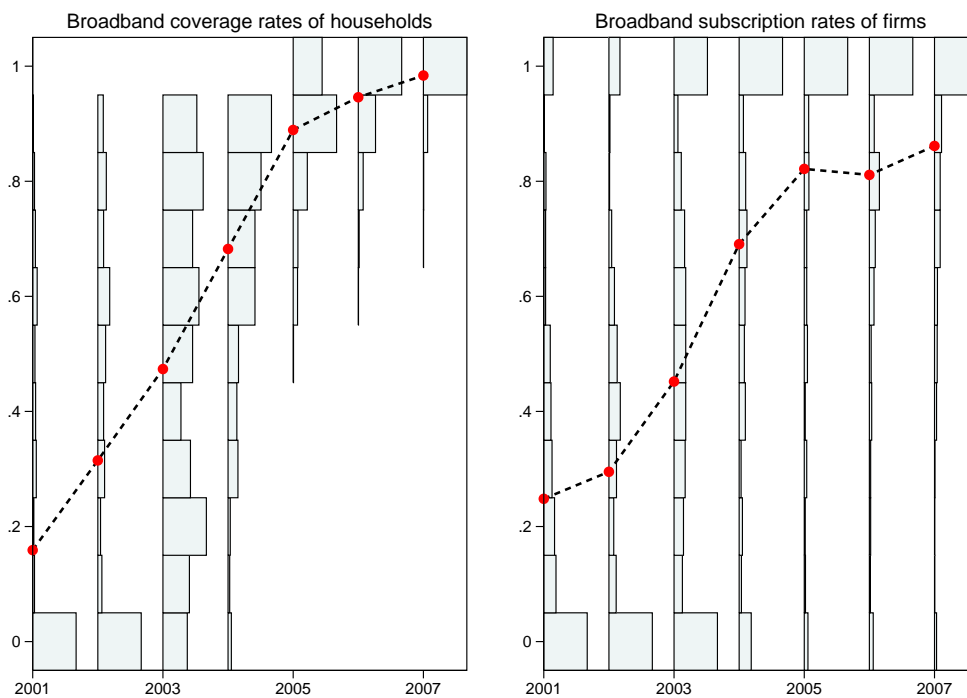
3.5.1 Identification of production functions

Randomizing broadband adoption is not feasible: We cannot in practice force firms to adopt a new technology. One can, however, think of a social experiment which randomizes broadband coverage at the municipality level. The randomization breaks the correlation between coverage rates and unobserved determinants of

¹⁴The Norwegian territory covers about 149,400 square miles, an area about the size of California or Germany, with around 13 % and 6 % of those regions' populations (in 2008), respectively. The country is dominated by mountainous or high terrain, as well as a rugged coastline stretching about 1,650 miles, broken by numerous fjords and thousands of islands.

¹⁵The reason is that the transmission of broadband signals through fiber-optic cables required installation of local access points. In areas with challenging topography and landscapes, it was more difficult and expensive to install the local access points and the fiber-optic cables. Furthermore, the existing infrastructure mattered for the marginal costs of installing cables to extend the coverage of broadband within a municipality and to neighboring areas.

Figure A1: Broadband coverage rates of households and broadband subscription rates of firms



Note: The graphs show the overall mean and distribution of broadband coverage rates of households and broadband subscription rates of firms across municipalities for each year during the period 2001–2007.

productivity. The intention of our IV-approach is to mimic this hypothetical experiment. Our source of exogenous variation in broadband adoption comes from the staged installation of broadband infrastructure, generating spatial and temporal variation in broadband coverage rates.

Below, we first describe how we use the roll-out of broadband coverage to address the endogeneity of internet adoption. Next, we discuss how we follow Levinsohn and Petrin (2003) in using intermediate inputs to control for correlation between input levels and the unobserved firm-specific productivity process.

IV model

Our IV strategy exploits that firms' access to and take-up of broadband internet changed as a side-effect of the expansion of broadband internet to households. This can be represented by the following system of equations, where the second stage is

$$y_{imkt} = x'_{imkt}\beta_0 + D_{imkt}x'_{imkt}\beta_1 + \eta_m + \tau_t + \lambda_k + \varepsilon_{imkt}, \quad (3.5)$$

while the first stages are given by

$$\begin{aligned}
 D_{imkt} &= x'_{imkt}\delta + Z_{mt}x'_{imkt}\phi + \gamma_m + \theta_t + \sigma_k + \nu_{imkt} \\
 D_{imkt}x_{1,imkt} &= x'_{imkt}\delta_1 + Z_{mk}x'_{imkt}\phi_1 + \gamma_{1,m} + \theta_{1,t} + \sigma_{1,k} + \nu_{1,imkt} \\
 &\vdots = \vdots \\
 D_{imkt}x_{n,imkt} &= x'_{imkt}\delta_n + Z_{mk}x'_{imkt}\phi_n + \gamma_{n,m} + \theta_{n,t} + \sigma_{n,k} + \nu_{n,imkt}
 \end{aligned} \tag{3.6}$$

where y is (log) value-added, D is a dummy variable for whether the firm uses broadband internet, Z is the fraction of households in a municipality with access to broadband internet, and x is a vector of n inputs (in log) and a constant term. The vector of parameters β_0 consists of output elasticities for the firms that have not adopted broadband internet, and β_1 is the difference in the output elasticities between the firms that have adopted broadband internet and the ones that have not. In all cases, subscript i denotes firm, subscript k denotes industry, subscript m denotes municipality and subscript t denotes year. All equations include a full set of municipality indicators (η_m, γ_m), year indicators (τ_t, θ_t), and (4 digit) industry indicators (σ_k, λ_k). Throughout the paper, the standard errors are clustered at the municipality level and robust to heteroskedasticity.

Unobservable determinants of production that are fixed at the municipality or industry level will be controlled for through the municipality and industry indicators, just like common time shocks are absorbed by the year indicators. To check that the estimated output elasticities are not confounded by time-varying observable factors, we report results with and without a large set of time-varying controls. In particular, we show that our estimates change little if we add control variables for the potential supply and demand factors discussed above.

Assessing the IV model

Our IV model – which controls for municipality (time and industry) fixed effects – is motivated by two features of the program. First, most of the supply and demand factors tend to vary little over time. Second, the timing of the roll-out is unlikely to co-vary with key correlates of production.

To investigate whether the data is consistent with these program features, we first regress Z_{mt} on municipality, time and industry fixed effects as well as time-varying supply and demand factors. We find that 89 % of the variation in broadband coverage can be attributed to time-invariant municipality and industry characteristics and common time effects, while less than 1 % of the variation in broadband coverage can be attributed to a large set of time-varying variables.¹⁶

Second, we examine the relationship between the timing of broadband roll-out and baseline municipality characteristics. To this end, we estimate the following equation

$$\Delta Z_{mt} = \eta_m + [\theta_t \times W_{m,2000}]' \psi_t + \epsilon_{mt} \tag{3.7}$$

where $\Delta Z_{mt} = Z_{mt} - Z_{m,t-1}$, θ_t is a vector of year fixed effects, and $W_{m,2000}$ includes municipality-level information from year 2000 on demography, average levels of inputs and output, industry structures and

¹⁶The time-varying variables include demographic factors (income level, education, share of population residing in a densely populated locality, size of population, and level of unemployment), inputs and output (municipality averages of revenues, intermediates, capital stock, number of workers and wage bill), industry structure (number of firms, employment share in manufacturing, employment share in wholesale, and employment share in services) and skill composition (shares of wages and workers by skill level).

skill composition. Demography includes income level, education, share of population residing in a densely populated locality (an urbanization indicator), size of population and level of unemployment. For inputs and output we have included municipality averages of revenues, intermediates, capital stock, number of workers and wage bill. As measures of industry structure, we use number of firms, employment share in manufacturing, employment share in wholesale, and employment share in services. The last group of municipality characteristics is the skill composition, measured as the shares of wages and workers by skill level.

Appendix Figure A5 plots the estimated coefficients from the vector ψ_t for every t (and the associated 95 % confidence intervals). Our results indicate that broadband expansion is positively related to urbanization until 2002. From 2003 onwards, there appears to be no systematic relationship between the timing of the broadband expansion and this variable. But more importantly, the timing of the expansion does not correlate with baseline industry structure, the levels of output and inputs, as well as the skill composition. Nevertheless, in the robustness analysis we demonstrate that our estimates change little when allowing for different underlying time trends in firm productivity or labor outcomes across municipalities.

Extension of the IV model

The IV approach breaks the correlation between D_{imkt} and ε_{imkt} . However, it is still possible that the factor inputs x_{imkt} are correlated with broadband adoption and unobserved productivity, which could bias the IV estimates of equation (3.5). Following Levinsohn and Petrin (2003, hereafter LP), we take a more structural approach to address this threat to identification of the production function.¹⁷

LP use a structural model of an optimizing firm to derive the conditions under which intermediate inputs can be used to proxy for unobserved productivity in the production function. The error term ε_{imkt} is assumed to be additively separable in a transmitted component (ω_{imkt}) and an i.i.d. component (χ_{imkt}). The key difference between ω_{imkt} and χ_{imkt} is that the former is a state variable, and therefore impacts the firm's decision rule, while the latter has no impact on the firm's decision. The intermediate input demand function depends on the firm-specific state variables, ω_{imkt} and k_{imkt} ,

$$a_{imkt} = g_t(\omega_{imkt}, k_{imkt}). \quad (3.8)$$

and it must be monotonic in ω for all relevant k .¹⁸ The monotonicity condition for intermediate inputs means that conditional on capital, profit maximizing behavior must lead more productive firms to use more intermediate inputs.

The monotonicity allows $g_t(\omega_{imkt}, k_{imkt})$ to be inverted to yield ω as a function of intermediate inputs and capital, $\omega_{imkt} = \omega_t(a_{imkt}, k_{imkt})$. By expressing the unobserved productivity variable ω_{imkt} as a function of observables, we are able to control for ω_{imkt} in the second stage equation:

$$y_{imkt} = x'_{imkt}\beta_0 + D_{imkt}x'_{imkt}\beta_1 + \eta_m + \tau_t + \lambda_k + \omega_t(a_{imkt}, k_{imkt}) + \chi_{imkt}. \quad (3.9)$$

As in Olley and Pakes (1996) and LP, we use a polynomial expansion in a and k to approximate $\omega_t(\cdot)$. By

¹⁷Levinsohn and Petrin (2003) extend on Olley and Pakes (1996) by using intermediate inputs instead of investments as a proxy for unobserved productivity. This addresses the problem that investment is zero in a non-trivial number of cases. Both studies also show how to address endogenous exit and entry of firms. Because our information on broadband adoption comes from repeated cross-sections (rather than panel data), we abstract from endogenous entry and exit.

¹⁸For simplicity, we assume as Olley and Pakes (1996) and Levinsohn and Petrin (2003) that capital is the only state variable over which the firm has control, while intermediates, labor and broadband internet are viewed as non-dynamic input factors.

simultaneous estimation of the first stage equations in (3.6) and the second stage equation in (3.9), we obtain consistent estimates of β_{u0} , β_{s0} , β_{k1} , β_{u1} , β_{s1} , and $\Phi_t(a_{imkt}, k_{imkt}) = \beta_{k0}k_{imkt} + \omega_t(a_{imkt}, k_{imkt})$.

While these output elasticities are sufficient to assess how broadband adoption affects labor productivity, we need to identify β_{k0} to recover the full shift in production technology. Because k_{imkt} is collinear with the non-parametric function $\omega_t(a_{imkt}, k_{imkt})$, further assumptions are necessary.¹⁹ Details on how β_{k0} is estimated are found in Appendix B.

We also apply the LP method when estimating the standard production function where any output driven by broadband internet will show up as TFP. Specifically, we express the unobserved productivity variable ω_{imkt} as a function of observables, $\omega_{imkt} = \omega_t(a_{imkt}, k_{imkt})$, and estimate the following regression model:

$$y_{imkt} = x'_{imkt}\beta + \rho_m + \varphi_t + \kappa_k + \omega_t(a_{imkt}, k_{imkt}) + \pi_{imkt}. \quad (3.10)$$

By using a polynomial expansion in a and k to approximate $\omega_t(\cdot)$ we obtain consistent estimates of β_u , β_s and $\Phi_t(a_{imkt}, k_{imkt}) = \beta_k k_{imkt} + \omega_t(a_{imkt}, k_{imkt})$. As in the case where firms can change their technology by adopting broadband internet, we impose further restrictions in order to identify β_k (see Appendix B).

3.5.2 Estimated production functions

Table A2: Production function estimates without broadband internet

	OLS	LP
<i>Panel B: 1 skill category</i>		
Capital	0.0891*** (0.00474)	0.1888*** (0.018)
Labor	0.833*** (0.00741)	0.650*** (0.00991)
<i>Panel B: 2 skill categories</i>		
Capital	0.110*** (0.00555)	0.225*** (0.019)
Unskilled	0.301*** (0.00898)	0.225*** (0.00703)
Skilled	0.423*** (0.0155)	0.303*** (0.0149)
Observations	18,675	18,675

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: The survey sample consists of joint-stock firms recorded in the internet survey over the period 2001-2007. Sampling weights are used to ensure representative results for the population of joint-stock firms. (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. All regressions include fixed effect for year, municipality and industry.

Table A2 reports estimates from a standard Cobb-Douglas production function without broadband adoption. The first column report OLS estimates, while the last column use the method for estimating production functions proposed by Levinsohn and Petrin (2003). We see that our estimates align well with the findings in previous studies. As predicted by theory, OLS overstates the labor coefficients because the level of inputs chosen is positively correlated with unobserved productivity. The magnitudes of the output elasticities of

¹⁹ β_{k1} is identified as the interaction of capital with D_{imkt} provides independent variation. Note also that the intercept in the production function is not separately identified from the mean of $E[\omega_{imkt}|\omega_{imkt-1}]$ without some further restriction.

Table A3: Production function estimates with broadband internet

	Output elasticities		
	OLS	IV	LP
Capital	0.108*** (0.0069)	0.113*** (0.0170)	0.232*** (0.032)
Unskilled	0.302*** (0.0131)	0.322*** (0.0248)	0.256*** (0.0214)
Skilled	0.335*** (0.0155)	0.310*** (0.0253)	0.231*** (0.0196)
Internet \times Capital	-0.0016 (0.0086)	-0.0074 (0.0248)	-0.0070 (0.0202)
Internet \times Unskilled	-0.0198 (0.0145)	-0.0553 (0.0396)	-0.0574* (0.0333)
Internet \times Skilled	0.136*** (0.0143)	0.202*** (0.0433)	0.141*** (0.0373)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Note: Baseline firm sample over the period 2001-2007, consisting of 18,675 joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. All regressions include fixed effect for year, municipality and industry. To test for strong skill bias, we use the method of Holm (1979).

capital and labor are comparable to what found in previous studies using micro data (see e.g. Pavcnik (2002); Fox and Smeets (2011)).

Table A3 reports production function estimates with broadband internet. The full set of first stage results are reported in Table A4. The first stages are strong, with large F-statistics on the excluded instruments which mean weak instrument bias is not a concern. Table A3 reports production function estimates based on OLS, IV and LP. In all cases, we find that broadband adoption in firms favors skilled labor by increasing its relative productivity. Indeed, the output elasticity of skilled labor increases substantially, while broadband internet seems to be a substitute for unskilled labor. By comparison, there is little if any change in the output elasticity of capital. We refer to Akerman, Gaarder and Mogstad (2013) for a further discussion of the technological change in production induced by broadband adoption.

Table A4: First stage regressions: 2 skill categories.

	Dependent variable:	Internet	Internet × Capital	Internet × Unskilled	Internet × Skilled
		(1)	(2)	(3)	(4)
<i>Panel A:</i> IV	Capital	0.0237*** (0.00810)	0.223*** (0.0454)	0.294*** (0.0937)	0.283*** (0.0949)
	Unskilled	0.0341*** (0.0122)	0.236*** (0.0585)	0.527*** (0.134)	0.485*** (0.146)
	Skilled	0.0596*** (0.0136)	0.392*** (0.0647)	0.784*** (0.160)	0.877*** (0.168)
	Coverage	0.589*** (0.167)	1.709** (0.806)	-0.270 (1.964)	0.393 (1.987)
	Coverage × Capital	-0.0159* (0.00928)	0.609*** (0.0500)	-0.192* (0.109)	-0.176 (0.110)
	Coverage × Unskilled	-0.0294** (0.0150)	-0.167** (0.0690)	0.317* (0.162)	-0.355** (0.177)
	Coverage × Skilled	0.00536 (0.0171)	-0.114 (0.0745)	0.0206 (0.198)	0.620*** (0.202)
	F-value (instruments)	18.3	48.8	14.3	13.0
<i>Panel B:</i> LP	Capital	0.0589** (0.0247)	-0.227** (0.113)	0.551* (0.282)	0.504* (0.295)
	Unskilled	0.0251** (0.0122)	0.169*** (0.0575)	0.409*** (0.133)	0.352** (0.146)
	Skilled	0.0493*** (0.0136)	0.310*** (0.0645)	0.643*** (0.159)	0.717*** (0.166)
	Coverage	0.596*** (0.166)	1.936** (0.790)	0.0365 (1.955)	0.824 (1.976)
	Coverage × Capital	-0.0143 (0.00910)	0.602*** (0.0479)	-0.182* (0.106)	-0.168 (0.107)
	Coverage × Unskilled	-0.0277* (0.0150)	-0.170** (0.0669)	0.329** (0.161)	-0.344* (0.176)
	Coverage × Skilled	0.00256 (0.0169)	-0.128* (0.0734)	-0.0204 (0.196)	0.571*** (0.199)
	F-value (instruments)	18.8	51.8	15.0	13.5

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: Baseline firm sample over the period 2001-2007, consisting of 18,675 joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. Panel A (B) provides the first stage results for the IV (LP) method.

In Table A5, we challenge the validity of the empirical strategy in several ways. Column 1 repeats the baseline estimates, while columns 2 and 3 include a wide range of controls for time-varying demographic and industry characteristics. When we include these covariates, we find that the estimated output elasticities are quite similar to our baseline estimates. An important requirement for our instrument to be valid is that the expansion of broadband internet is unrelated to different underlying time trends in firm productivity across municipalities. Columns 4-6 of Table A5 address this threat to identification. We first interact baseline (year 2000) covariates with a linear (column 4) and quadratic time trend (column 5). By including these controls in both the first and second stages, we allow the expansion of broadband internet to be related to different underlying time trends in productivity across municipalities, depending on their initial characteristics. A limitation of these robustness checks is that baseline characteristics may not adequately capture differential

time trends across municipalities. Column 6 includes linear municipality-specific time trends in both the first and second stages. It is reassuring to find that neither the IV estimates nor the LP estimates change significantly when allowing for differential time trends across municipalities.

Table A5: Specification checks for production function

	Time interacted with covariates						
	Baseline (1)	Covariates (2)	(3)	Linear (4)	Quadratic (5)	Linear municipal trends (6)	
<i>Panel A:</i>							
IV	Internet × Capital	-0.00741 (0.0248)	-0.0104 (0.0253)	-0.0132 (0.0248)	-0.0183 (0.0275)	-0.0187 (0.0265)	-0.00101 (0.0236)
	Internet × Unskilled	-0.0553 (0.0396)	-0.0600 (0.0406)	-0.0598 (0.0403)	-0.0677 (0.0422)	-0.0664 (0.0420)	-0.0305 (0.0431)
	Internet × Skilled	0.202*** (0.0433)	0.202*** (0.0441)	0.192*** (0.0438)	0.195*** (0.0451)	0.190*** (0.0445)	0.163*** (0.0410)
<i>Panel B:</i>							
LP	Internet × Capital	-0.00701 (0.0202)	-0.00950 (0.0205)	-0.0113 (0.0199)	-0.0188 (0.0216)	-0.0184 (0.0209)	-0.00247 (0.0202)
	Internet × Unskilled	-0.0574* (0.0333)	-0.0608* (0.0340)	-0.0624* (0.0339)	-0.0707** (0.0352)	-0.0692** (0.0348)	-0.0412 (0.0334)
	Internet × Skilled	0.141*** (0.0373)	0.140*** (0.0377)	0.132*** (0.0370)	0.134*** (0.0375)	0.129*** (0.0373)	0.105*** (0.0369)
Controls	Demographic	✓	✓	✓	✓	✓	✓
	Industry						

* p < 0.10, ** < 0.05, *** p < 0.01.

Note: Baseline firm sample over the period 2001-2007, consisting of 18,675 joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. (Un)Skilled comprises workers with(out) high school diploma. The standard errors are clustered at the municipality level and robust to heteroskedasticity. All regressions include fixed effect for year, municipality and industry. Column (2) adds demographic controls to the baseline model, including municipality-level information on average household income, mean years of schooling, share of population residing in a densely populated locality, size of population and level of unemployment. Column (3) also includes industry controls, consisting of municipality averages of revenues, intermediates, capital stock, number of workers and wage bills as well as employment share in manufacturing, employment share in wholesale, employment share in services, and shares of wages and workers by skill level. Columns (4) and (5) interact linear and quadratic time trends with baseline values (year 2000) of these covariates. Column (6) includes municipality-specific linear time trends.

3.6 Implications of broadband internet for TFP

3.6.1 Dispersion of TFP across firms

To assess the quantitative importance of productivity-influencing factors, Syverson (2011) suggest comparing the dispersion in TFP estimates from production functions with and without these factors. We examine the contribution from broadband internet by comparing the results from the standard production function equation

$$y_{imkt} = x'_{imkt}\beta + \rho_m + \varphi_t + \kappa_k + \nu_{imkt} \quad (3.11)$$

to those produced by equation (3.5) where we allow broadband internet to shift the production technology. We define

$$\begin{aligned} A_{it}^0 &\equiv e^{\nu_{imkt}} \\ A_{it}^1 &\equiv e^{\hat{\epsilon}_{imkt}} \end{aligned}$$

as measures of the TFP of firms in the specification with (A_{it}^1) and without (A_{it}^0) broadband adoption in the production function.

To measure the extent to which broadband internet explains firm productivity, we compare the dispersion of A_{it}^0 and A_{it}^1 across firms in our baseline sample. The results are reported in Panel A of Table A6. Given that equation (3.5) nests equation (3.11), we expect less dispersion in A_{it}^1 as compared to A_{it}^0 . As expected, the dispersion in TFP across firms is reduced by including broadband internet in the production function estimation. Using the IV and LP estimates, the standard deviation in TFP declines by 2.4 percent and 3.3 percent, respectively. The results for the ratios between different percentiles suggest the upper part of the TFP distribution is more influenced by taking into account broadband internet.

3.6.2 Aggregate TFP

We now shift attention to the implications for aggregate TFP in the economy, which is a more welfare relevant measure than TFP at the firm level. We follow Olley and Pakes (1996) in defining aggregate TFP, A_t , as

$$A_t \equiv \sum_i s_{it} A_{it}$$

where s_{it} is defined as the market share (the share of value added) of firm i in year t . Olley and Pakes (1996) show how aggregate TFP in an economy can be decomposed into an unweighted average of firm-level TFP and the covariance between market share and firm-level TFP. Formally, this can be written as

$$A_t = \bar{A}_t + \sum_i \Delta s_{it} \Delta A_{it} \quad (3.12)$$

where \bar{A}_t indicates the unweighted average of A_{it} , $\Delta s_{it} \equiv s_{it} - \bar{s}_{it}$ and $\Delta A_{it} \equiv A_{it} - \bar{A}_{it}$. The first term on the right hand side therefore captures the unweighted average of TFP across firms and the second term the covariance between market share and TFP. The covariance term shows the extent to which production (and production factors) is allocated to firms with higher unobserved productivity.

We use the measures of A_{it}^1 and A_{it}^0 to perform this decomposition. In Panel B of Table A6, we report the share of aggregate TFP that can be attributed to the covariance term ($\frac{\sum_i \Delta s_{it} \Delta A_{it}}{A_t}$). Using the IV and LP

Table A6: Total factor productivity, variance and decomposition.

Broadband internet:	OLS		IV		LP	
	Without	With	Without	With	Without	With
<i>Panel A: Variation in TFP</i>						
Standard deviation	0.84	0.77	0.84	0.82	1.21	1.17
90/50 percentiles	1.73	1.69	1.73	1.70	2.03	1.99
50/10 percentiles	1.61	1.59	1.61	1.60	1.96	1.94
<i>Panel B: Olley-Pakes decomposition</i>						
Share of covariance term in aggregate TFP	0.41	0.37	0.41	0.39	0.52	0.49
Total aggregate productivity	2.14	1.97	2.14	2.06	4.11	3.71
Unweighted average term	1.26	1.24	1.26	1.25	1.97	1.89
Covariance term	0.87	0.73	0.87	0.81	2.14	1.83

Note: The table uses the survey sample of 17,857 joint-stock firms for which we observe broadband internet adoption. Panel A displays various measures of variation for the residual term (total factor productivity) from two different production functions. The first (in columns marked “Without”) reports the case when broadband internet adoption is not accounted for as described in equation (3.11). The second (in columns marked “With”) reports the case when broadband internet adoption is accounted for as described in equation (3.5). Panel B calculates what share of the aggregate productivity in the sample can be attributed to the covariance term in the Olley-Pakes decomposition as reported in equation (3.12).

estimates, the contribution of the covariance term is reduced by 1.7–2.8 percentage points when we include broadband internet in the production function estimation. This suggests that broadband internet accounts for some of the allocation of production to more productive firms.

3.7 Conclusion

In 2009, the U.S. Congress asked for a plan that would provide affordable broadband service to all America’s citizens. In other OECD countries, there has been similar calls for public funding to roll out broadband infrastructure. While government agencies are projecting broadband penetration rates to be important for productivity, there is little scientific evidence to substantiate these claims.

We contribute by examining how the adoption of new broadband internet induces productivity changes and explains dispersion in TFP across firms. Our analysis employs rich Norwegian data with firm-level information on value added, capital, labor by skill level, and broadband adoption. We begin by estimating a standard production function where any output driven by broadband internet will show up as TFP. Next, we estimate a production function where firms can change their technology by adopting broadband internet. We allow for the possibility that broadband adoption is factor-neutral or factor-biased, and therefore, directly or indirectly affect output. A public program with limited funding rolled out broadband access points, and provides plausibly exogenous variation of broadband adoption in firms.


We find that broadband internet accounts for about 2 to 3 percent of the standard deviation in TFP across firms. We decompose aggregate TFP in the economy into an unweighted average of firm-level TFP and the covariance between market share and firm-level TFP. We find that the contribution of the covariance term to aggregate TFP is reduced substantially when we include broadband internet in the production function estimation. This suggests that broadband internet accounts for some of the allocation of production to more

productive firms.

Taken together, these findings have important implications for the debate about what explains productivity differences. Our results are consistent with previous studies in pointing to the importance of factors that are internal to the firm. However, this does not imply that there is limited scope for policy to influence productivity. On the contrary, our estimates suggest that policy increasing the availability of broadband internet induces differences in firms' choice of technology, which can be important in explaining the dispersion of TFP. This illustrates that even if firms control the key factors of productivity, the external products of the operating environment are important in influencing the choices firms make.

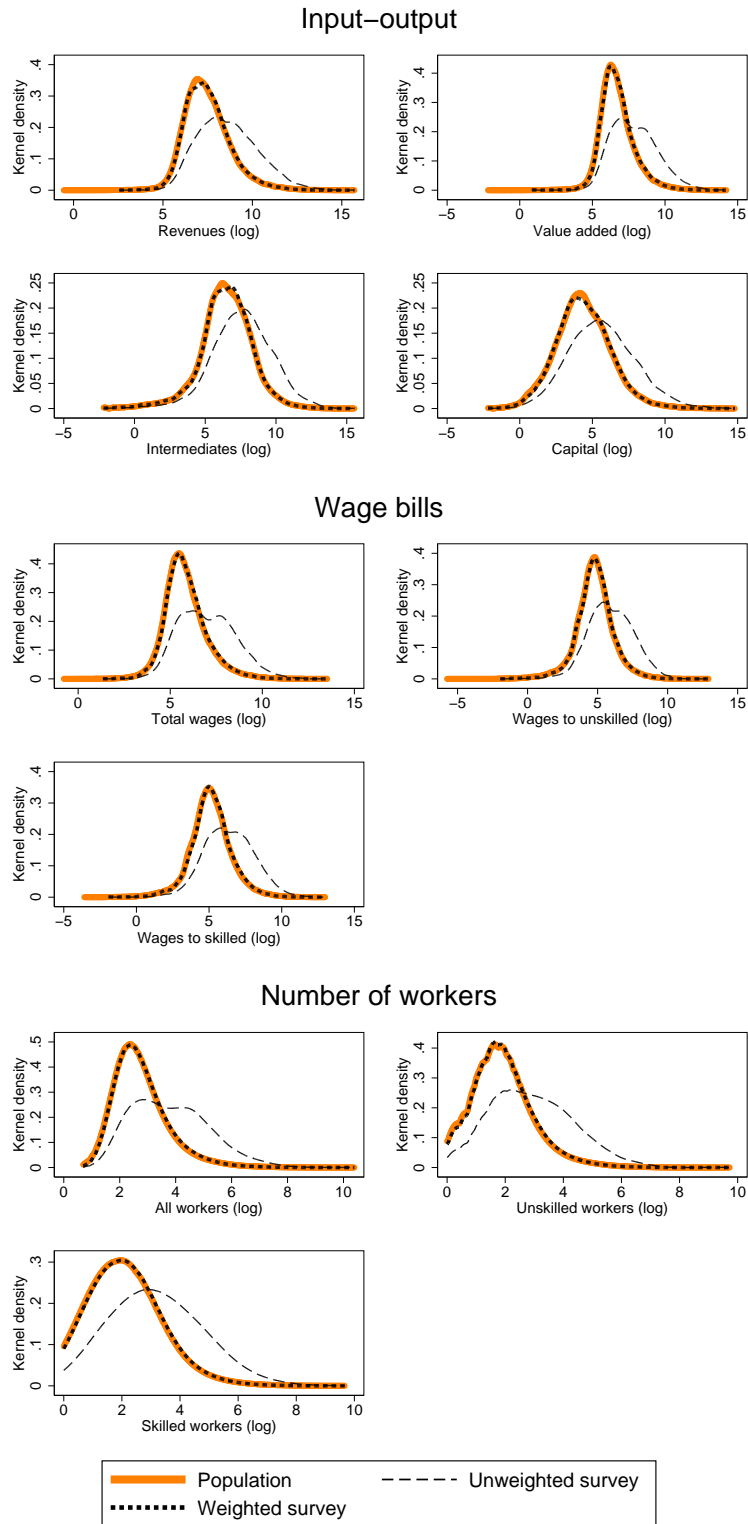
3.A Appendix A: Additional Tables and Figures

Figure A1: Distribution of firms by industry



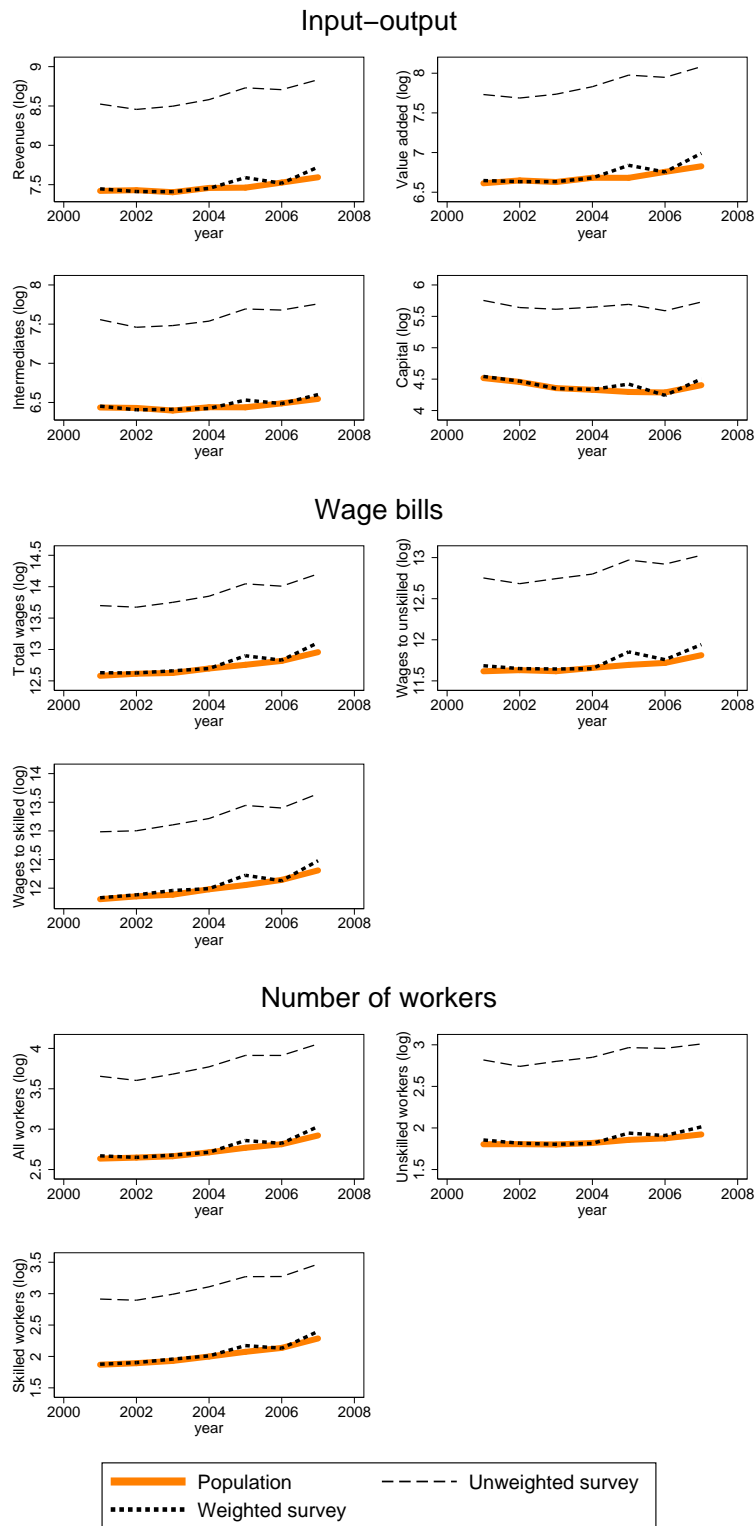
Note: The survey sample consists of joint-stock firms recorded in the internet survey. Sampling weights are used to ensure representative results for the population of joint-stock firms. We exclude firms that are carrying out extraction of natural resources (including oil, gas and fish).

Figure A2: Cross-sectional distribution of key firm variables




Note: The survey sample consists of joint-stock firms recorded in the internet survey over the period 2001-2007. The weighted survey uses sampling weights to ensure representative results for the population of joint-stock firms. We exclude firms that are carrying out extraction of natural resources (including oil, gas and fish). Detailed descriptions of the variables are given in Appendix Table A1.

Figure A3: Time trends in key firm variables




Note: The survey sample consists of joint-stock firms recorded in the internet survey over the period 2001-2007. The weighted survey uses sampling weights to ensure representative results for the population of joint-stock firms. We exclude firms that are carrying out extraction of natural resources (including oil, gas and fish). Detailed descriptions of the variables are given in Appendix Table A1

Figure A4: Geographical distribution of broadband coverage rates.



Note: The graphs show the geographical distribution of broadband coverage rates of households in 2001, 2003 and 2005.

Figure A5: Timing of broadband expansion and baseline covariates



Note: We report estimates from equation (2.5) of the vector ψ_t for every t (and the associated 95 % confidence intervals). Demography includes income level, education, share of population residing in a densely populated locality (an urbanization indicator), size of population and level of unemployment. As measures of industry structure, we use number of firms, employment share in manufacturing, employment share in wholesale, and employment share in services. For inputs and output we have included municipality averages of revenues, intermediates, capital stock, number of workers and wage bill.

Table A1: Variable definitions

Variable	Description
<i>Firm accounts</i>	Source: The Account Statistics.
Revenues	Total sales by a firm in year t .
Intermediates	Procurement of materials and intermediate inputs of a firm in year t .
Capital	Value of total fixed assets of a firm in year t .
Value added	Sales minus intermediates of a firm in year t .
Industry	4-digit code classifying a firm's main activity in year t according to the Standard Industry Classification (SIC2002) system.
Municipality	4-digit code for the municipality in which a firm is located in year t .
<i>Internet variables</i>	Source: The community survey on ICT in firms
Broadband	Dummy variable for whether a firm has adopted broadband internet (speed at or above 256 kilobits per second) in year t .
<i>Employees</i>	Source: Register of Employers and Employees and the Wage Statistics Survey.
Annual wages	Annual pre-tax wages in year t
<i>Internet coverage</i>	Source: Norwegian Ministry of Government Administration.
Coverage rate	Fraction of households in a given municipality with access to broadband internet (speed at or above 256 kilobits per second) in year t .
<i>Demographic controls</i>	Source: Central Population Register.
Urbanization	Population share living in densely populated area in a given municipality in year t .
Income	Average annual disposable income across individuals aged 16–59 years in a given municipality in year t .
Education	Average years of schooling across individuals aged 16–59 in a given municipality in year t .
Unemployment	Unemployment rate among individuals aged 16–59 in a given municipality in year t .
<i>Industry and firm controls</i>	Source: The Account Statistics and Register of Employers and Employees.
Share of skilled workers	Share of employed workers with high school diploma in a given municipality in year t .
Share of total wages to skilled workers	Share of the total wage bill paid to workers with high school diploma in a given municipality in year t .
Share of employment by industry	Share of workers in the manufacturing/wholesale/service industry in a given municipality in year t .
Average input levels	Average level of capital stock/value added/number of workers/wages paid/revenues across firms in a given municipality in year t .

3.B Appendix B: Estimation details

By simultaneous estimation of equations (2.4) and (2.7), we obtain consistent estimates of β_{u0} , β_{s0} , β_{k1} , β_{u1} , β_{s1} , and $\Phi_t(a_{imkt}, k_{imkt}) = \beta_{k0}k_{imkt} + \omega_t(a_{imkt}, k_{imkt})$. While these output elasticities are sufficient to assess how broadband adoption affects labor productivity, we need to identify β_{k0} to recover the full shift in production technology. Because k_{imkt} is colinear with the non-parametric function $\omega_t(a_{imkt}, k_{imkt})$, further assumptions are necessary.

Assuming that ω_{imkt} follows a first-order Markow process, we can write

$$\omega_{imkt} = E[\omega_{imkt}|\omega_{imkt-1}] + \xi_{imkt}.$$

This simply decomposes ω_{imkt} into its conditional expectation at time $t-1$, $E[\omega_{imkt}|\omega_{imkt-1}]$, and a deviation from that expectation, ξ_{imkt} . If the capital stock is pre-determined and current investment (which will react to productivity shocks) takes one period before it comes productive, it follows that

$$E[\xi_{imkt}k_{imkt}] = 0.$$

This is the moment which LP use to identify the capital coefficient. Roughly speaking, variation in k_{imkt} conditional on ω_{imkt-1} is the exogenous variation used for identification. To operationalize this approach in a GMM context, note that given a guess on the capital coefficient β_{k0} , we can rewrite unobserved productivity as

$$\omega_{imkt}(\beta_{k0}) = \hat{\Phi}_{imkt} - \beta_{k0}k_{imkt}.$$

Given these $\omega_{imkt}(\beta_{k0})$, we compute ξ_{imkt} by non-parametrically regressing $\omega_{imkt}(\beta_{k0})$'s on $\omega_{imkt-1}(\beta_{k0})$'s and a constant term; we then form the residual

$$\xi_{imkt}(\beta_{k0}) = \omega_{imkt}(\beta_{k0}) - \hat{\Psi}(\omega_{imkt-1}(\beta_{k0}))$$

where $\hat{\Psi}(\omega_{imkt-1}(\beta_{k0}))$ are predicted values from the non-parametric regression.

The $\xi_{imkt}(\beta_{k0})$'s are used to form a sample analogue to the above moment. i.e.

$$\frac{1}{T} \frac{1}{N} \sum_t \sum_i \xi_{imkt}(\beta_{k0}) \cdot k_{imkt}.$$

We estimate β_{k0} by minimizing the GMM criterion function

$$Q(\beta_{k0}) = \min_{\beta_{k0}} \left(\frac{1}{N} \frac{1}{T_{i1}} \sum_i \sum_{t=T_{i0}}^{T_{i1}} \xi_{imkt}(\beta_{k0}) \cdot k_{imkt} \right)^2$$

where i indexing firms and T_{i0} and T_{i1} index the second and last period in which firm i is observed.

Because our baseline sample is a repeated cross-section (rather than panel data), we adjust the above estimation procedure. Exploiting the random sampling of firms, we can identify β_{k0} from the moment

$$E[\bar{\xi}_{mkt}, \bar{k}_{mkt}] = 0.$$

where the municipality average of a variable is denoted by upper bar. By applying the above procedure to

our panel data at the municipality level, we obtain the GMM criterion function

$$Q(\beta_{k0}) = \min_{\beta_{k0}} \left(\frac{1}{M} \frac{1}{T_{m1}} \sum_m \sum_{t=T_{m0}}^{T_{m1}} \sqrt{N_{mkt}} \bar{\xi}_{mkt}(\beta_{k0}) \cdot \bar{k}_{mkt} \right)^2$$

where T_{m0} and T_{m1} index the second and last period municipality m is observed and N_{mkt} is the number of firms in municipality m in period t . To obtain standard errors on β_{k0} in Table 2.2, we use bootstrap while clustering by municipality.

Bibliography

- AABERGE, R. (1997): “Interpretation of Changes in Rank-Dependent Measures of Inequality,” *Economics Letters*, 55(2), 215–219.
- ABOWD, J. M., J. HALTIWANGER, R. JARMIN, J. LANE, P. LENGERMANN, K. MCCUE, K. MCKINNEY, AND K. SANDUSKY (2005): “The relation among human capital, productivity, and market value: Building up from micro evidence,” in *Measuring capital in the new economy*, pp. 153–204. University of Chicago Press.
- ACEMOGLU, D. (2003): “Patterns of Skill Premia,” *Review of Economic Studies*, 70(2), 199–230.
- ACEMOGLU, D., AND D. H. AUTOR (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in *Handbook of Labor Economics Vol. 4*, ed. by O. Ashenfelter, and D. E. Card, chap. 12, pp. 1043–1171. Elsevier, Amsterdam.
- AKERMAN, A., I. GAARDER, AND M. MOGSTAD (2013): “The skill complementarity of broadband internet,” Discussion paper, IZA Discussion Paper No. 7762.
- ANDERSON, S. P., A. DE PALMA, AND B. KREIDER (2001): “Tax Incidence in Differentiated Product Oligopoly,” *Journal of Public Economics*, 81(2), 173–192.
- ATKINSON, A. B., L. RAINWATER, AND T. M. SMEEDING (1995): *Income Distribution in OECD Countries: Evidence from the Luxembourg Income Study*. Organisation for Economic Co-operation and Development Paris.
- ATKINSON, T. (2013): “The “Mirrlees Review” of UK Taxation,” *Royal Economic Society. Newsletter*, 160, 5–6.
- AUTOR, D. H., AND D. DORN (2009): “This Job is ‘Getting Old’: Measuring Changes in Job Opportunities using Occupational Age Structure,” *American Economic Review: Papers & Proceedings*, 99(2), 45–51.
- (2013): “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103(5), 1553–1597.
- AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): “The Polarization of the U.S. Labor Market,” *American Economic Review Papers and Proceedings*, 96(2), 189–194.
- (2008): “Trends in US Wage Inequality: Revising the Revisionists,” *Review of Economics and Statistics*, 90(2), 300–323.

- AUTOR, D. H., L. F. KATZ, AND A. B. KRUEGER (1998): “Computing Inequality: Have Computers Changed the Labor Market?,” *Quarterly Journal of Economics*, 113(4), 1169–1213.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): “The Skill Content of Recent Technological Change: An Empirical Exploration,” *Quarterly Journal of Economics*, 118(4), 1279–1333.
- BANKS, J., R. BLUNDELL, AND A. LEWBEL (1996): “Tax Reform and Welfare Measurement: Do We Need Demand System Estimation?,” *The Economic Journal*, 106(438), 1227–1241.
- (1997): “Quadratic Engel Curves and Consumer Demand,” *Review of Economics and Statistics*, 79(4), 527–539.
- BEAUDRY, P., M. DOMS, AND E. LEWIS (2010): “Should the Personal Computer Be Considered a Technological Revolution? Evidence from U.S. Metropolitan Areas,” *Journal of Political Economy*, 118(5), 988 – 1036.
- BEAUDRY, P., AND D. A. GREEN (2003): “The Changing Structure of Wages in the US and Germany: What Explains the Differences?,” *American Economic Review*, 93(3), 573–602.
- (2005): “Changes in U.S. Wages, 1976–2000: Ongoing Skill Bias or Major Technological Change?,” *Journal of Labor Economics*, 23(3), 609–648.
- BERMAN, E. (2000): “Does Factor-Biased Technological Change Stifle International Convergence? Evidence from Manufacturing,” NBER Working Paper 7964.
- BERMAN, E., J. BOUND, AND Z. GRILICHES (1994): “Changes in the Demand for Skilled Labor within US Manufacturing: Evidence from the Annual Survey of Manufactures,” *Quarterly Journal of Economics*, 109(2), 367–397.
- BESLEY, T. J., AND H. S. ROSEN (1999): “Sales Taxes and Prices: An Empirical Analysis,” *National Tax Journal*, 52, 157–178.
- BHULLER, M., T. HAVNES, E. LEUVEN, AND M. MOGSTAD (2013): “Broadband internet: An information superhighway to sex crime?,” *The Review of Economic Studies*, 80(4), 1237–1266.
- BLACK, S. E., AND A. SPITZ-OENER (2010): “Explaining Women’s Success: Technological Change and the Skill Content of Women’s Work,” *Review of Economics and Statistics*, 92(1), 187–194.
- BLOOM, N., M. DRACA, AND J. VAN REENEN (2011): “Trade induced technical change? The impact of Chinese imports on innovation, IT and productivity,” Discussion paper, National Bureau of Economic Research.
- BLOOM, N., L. GARICANO, R. SADUN, AND J. VAN REENEN (2009): “The Distinct Effects of Information Technology and Communication Technology on Firm Organization,” NBER Working Paper 14975.
- BLOOM, N., R. SADUN, AND J. VAN REENEN (2012): “Americans Do IT Better: US Multinationals and the Productivity Miracle,” *American Economic Review*, 102(1), 167–201.
- BLOOM, N., M. SCHANKERMAN, AND J. VAN REENEN (2013): “Identifying technology spillovers and product market rivalry,” *Econometrica*, 81(4), 1347–1393.

- BLOOM, N., AND J. VAN REENEN (2007): "Measuring and explaining management practices across firms and countries," *The Quarterly Journal of Economics*, 122(4), 1351–1408.
- BLUNDELL, R., P. PASHARDES, AND G. WEBER (1993): "What Do We Learn About Demand Patterns from Micro Data," *The American Economic Review*, 83(3), 570–597.
- BOND, S., AND J. VAN REENEN (2007): "Microeconomic Models of Investment and Employment," in *Handbook of Econometrics*, ed. by J. J. Heckman, and E. E. Leamer, pp. 4417–4498. North Holland, London.
- BRESNAHAN, T. F., E. BRYNJOLFSSON, AND L. M. HITT (2002): "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-level evidence," *Quarterly Journal of Economics*, 117(1), 339–376.
- BROWNING, M., AND C. MEGHIR (1991): "The Effects of Male and Female Labor Supply on Commodity Demands," *Econometrica*, 59(4), 925–951.
- BRYNJOLFSSON, E., AND L. M. HITT (2003): "Computing Productivity: Firm-level Evidence," *Review of Economics and Statistics*, 85(4), 793–808.
- CARBONNIER, C. (2007): "Who Pays Sales Taxes? Evidence from French VAT reforms, 1987–1999," *Journal of Public Economics*, 91, 1219–1229.
- (2013): "The Incidence of Non-Linear Consumption Taxes," Discussion paper.
- CAROLI, E., AND J. VAN REENEN (2001): "Skill Biased Organizational Change? Evidence from a Panel of British and French Establishments," *Quarterly Journal of Economics*, 116(4), 1449–1492.
- CRAWFORD, I., M. KEEN, AND S. SMITH (2010): "Value Added Tax and Excises," in *Dimensions of Tax Design, The Mirrlees Review*, ed. by J. Mirrless, S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, and J. Poterba, chap. 4, pp. 275–422. Oxford University Press Inc., New York.
- CZERNICH, N., O. FALCK, T. KRETSCHMER, AND L. WOESSMANN (2011): "Broadband Infrastructure and Economic Growth," *Economic Journal*, 121(May), 505–532.
- DEATON, A., AND J. MUELLBAUER (1980): "An Almost Ideal Demand System," *The American Economic Review*, 70(3), 312–326.
- DECICCA, P., D. KENKEL, AND F. LIU (2013): "Who Pays Cigarette Taxes? The Impact of Consumer Price Search," *Review of Economics and Statistics*, 95(2), 516–529.
- DELIPALLA, S., AND M. KEEN (1992): "The Comparison between Ad Valorem and Specific Taxation under Imperfect Competition," *Journal of Public Economics*, 49, 351–367.
- DELIPALLA, S., AND O. O'DONNELL (2001): "Estimating Tax Incidence, Market Power and Market Conduct: the European Cigarette Industry," *International Journal of Industrial Organization*, 19(6), 885–908.
- DIÑARDO, J. E., AND J.-S. PISCHKE (1997): "The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?," *Quarterly Journal of Economics*, 112(1), 291–303.

- DOYLE, J. J., AND K. SAMPHANTHARAK (2008): “\$ 2.00 Gas! Studying the Effects of a Gas Tax Moratorium,” *Journal of Public Economics*, 92(3), 869–884.
- DRACA, M., R. SADUN, AND J. VAN REENEN (2007): “Productivity and ICTs: A Review of the Evidence,” in *Oxford Handbook of Information and Communication Technologies*, ed. by R. Mansell, C. Avgerou, D. Quah, and R. Silverstone, pp. 100–147. Oxford University Press.
- EINARSSON, A. (2007): “The Retail Sector in the Nordic Countries. A Comparative Analysis,” *Bifrist Journal of Social Science, Working Paper*, 1.
- FABINGER, M., AND E. G. WEYL (2012): “Pass-through and demand forms,” *This work is in progress. For a draft of notes, contact Glen Weyl at weyl@uchicago.edu.*
- FIRPO, S., N. M. FORTIN, AND T. LEMIEUX (2011): “Occupational Tasks and Changes in the Wage Structure,” IZA Discussion Paper 5542.
- FORMAN, C., A. GOLDFARB, AND S. GREENSTEIN (2012): “The Internet and Local Wages: A Puzzle,” *American Economic Review*, 102(1), 556–575.
- FOSTER, L., J. HALTIWANGER, AND C. SYVERSON (2008): “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?,” *American Economic Review*, 98(1), 394–425.
- FOX, J. T., AND V. SMEETS (2011): “Does Input Quality Drive Measured Differences In Firm Productivity?,” *International Economic Review*, 52(4), 961–989.
- GARICANO, L., AND P. HEATON (2010): “Information technology, organization, and productivity in the public sector: evidence from police departments,” *Journal of Labor Economics*, 28(1), 167–201.
- GOLDIN, C., AND L. F. KATZ (2007): “Long-Run Changes in the Wage Structure: Narrowing, Widening, Polarizing,” *Brookings Papers on Economic Activity*, 38(2), 135–168.
- GOOS, M., AND A. MANNING (2007): “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain,” *Review of Economics and Statistics*, 89(1), 118–133.
- HOLM, S. (1979): “A Simple Sequentially Rejective Multiple Test Procedure,” *Scandinavian Journal of Statistics*, 6(2), 65–70.
- HORTAÇSU, A., AND C. SYVERSON (2007): “Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices,” *Journal of Political Economy*, 115(2), 250–301.
- ICHNIEWSKI, C., AND K. SHAW (2003): “Beyond incentive pay: Insiders’ estimates of the value of complementary human resource management practices,” *Journal of Economic Perspectives*, pp. 155–180.
- IMBENS, G., AND K. KALYANARAMAN (2012): “Optimal Bandwidth Choice for the Regression Discontinuity Estimator,” *The Review of Economic Studies*, 79(3), 933–959.
- IMBENS, G. W., AND J. D. ANGRIST (1994): “Identification and Estimation of Local Average Treatment Effects,” *Econometrica*, 62(2), 467–475.
- JENSEN, J. B., AND L. G. KLETZER (2005): “Tradable Services: Understanding the Scope and Impact of Services Outsourcing,” Institute for International Economics Working Paper 05-9.

- JORGENSON, D. W., M. S. HO, AND K. J. STIROH (2008): “A Retrospective Look at the U.S. Productivity Growth Resurgence,” *Journal of Economic Perspectives*, 22(1), 3–24.
- KENKEL, D. S. (2005): “Are Alcohol Tax Hikes Fully Passed Through to Prices? Evidence from Alaska,” *American Economic Review*, pp. 273–277.
- KOSONEN, T. (2013): “More Haircut after VAT Cut? On the Efficiency of Service Sector Consumption Taxes,” .
- KRUEGER, A. B. (1993): “How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984-1989,” *Quarterly Journal of Economics*, 108(1), 33–60.
- LAZEAR, E. P. (2000): “Performance Pay and Productivity,” *American Economic Review*, 90(5), 1346–1361.
- LEE, D. S., AND T. LEMIEUX (2009): “Regression Discontinuity Designs in Economics,” Discussion paper, National Bureau of Economic Research.
- LEVINSOHN, J., AND A. PETRIN (2003): “Estimating Production Functions Using Inputs to Control for Unobservables,” *Review of Economic Studies*, 70(2), 317–341.
- MACHIN, S., AND J. VAN REENEN (1998): “Technology and Changes in Skill Structure: Evidence from Seven OECD Countries,” *Quarterly Journal of Economics*, 113(4), 1215–1244.
- MARION, J., AND E. MUEHLEGGGER (2011): “Fuel Tax Incidence and Supply Conditions,” *Journal of Public Economics*, 95(9), 1202–1212.
- MICHAELS, G., A. NATRAJ, AND J. VAN REENEN (forthcoming): “Has ICT Polarized Skill Demand? Evidence from Eleven Countries over 25 years,” *Review of Economics and Statistics*.
- OLINER, S. D., D. E. SICHEL, AND K. J. STIROH (2007): “Explaining a Productive Decade,” *Brookings Papers on Economic Activity*, 38(1), 81–152.
- OLLEY, S. G., AND A. PAKES (1996): “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica*, 64(6), 1263–1297.
- PAVCNIK, N. (2002): “Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean plants,” *Review of Economic Studies*, 69(1), 245–276.
- POTERBA, J. M. (1996): “Retail Price Reactions to Changes in State and Local Sales Taxes,” *National Tax Journal*, 49, 165–176.
- SCHMITZ JR, J. A. (2005): “What determines productivity? Lessons from the dramatic recovery of the US and Canadian iron ore industries following their early 1980s crisis,” *Journal of Political Economy*, 113(3), 582–625.
- SEADE, J. (1987): “Profitable Cost Increases and the Shifting of Taxation: Equilibrium Responses of Markets in Oligopoly,” .
- ST.MELD.NR. 38 (1997–1998): “IT-kompetanse i et regionalt perspektiv [A regional perspective on IT-competency],” Ministry of Trade and Industry.

- ST.MELD.NR. 49 (2003): “Breiband for kunnskap og vekst [Broadband for knowledge and growth],” Ministry of Trade and Industry.
- SYVERSON, C. (2004a): “Market Structure and Productivity: A Concrete Example,” *Journal of Political Economy*, 112(6), 1181–1222.
- (2004b): “Product Substitutability and Productivity Dispersion,” *The Review of Economics and Statistics*, 86(2), 534–550.
- (2011): “What Determines Productivity?,” *Journal of Economic Literature*, 49(2), 326–365.
- VAN BEVEREN, I. (2012): “Total Factor Productivity Estimation: A Practical Review,” *Journal of Economic Surveys*, 26(1), 98–128.
- YOUNG, D. J., AND A. BIELIŃSKA-KWAPISZ (2002): “Alcohol Taxes and Beverage Prices,” *National Tax Journal*, pp. 57–73.