High-speed broadband and academic achievement in teenagers: Evidence from Sweden

Erik Grenestam * Martin Nordin †

First draft: October 2016
This version: April 2019

Abstract

This study examines the effects of super-fast internet connections on the academic achievement of students in upper secondary school. We link detailed register data on around 250,000 students to local levels of access to optic fiber broadband to estimate the effect of broadband on student GPA using within-student variation in a difference-in-differences setup. We show that reaching full coverage in the student's parish of residence causes a GPA reduction of about 4 percent of a standard deviation in our preferred specification. The impact is greater for boys and students born to parents with a low level of education. The gender differential can explain part of a widening GPA-gap between boys and girls.

JEL classifications: J24, H52, I24, I28, O33
Keywords: Education, Broadband, Internet, High-school, GPA

*Department of Economics, Lund University. PO Box 730, SE-220 07, Lund, Sweden. E-mail: erik.grenestam@nek.lu.se
†AgriFood Economics Centre and Department of Economics, Lund University. PO Box 730, SE-220 07, Lund, Sweden.
The authors would like to thank Kaveh Majlesi, Petter Lundborg, Alessandro Martinello and Jan Bietenbeck for their feedback as well as seminar participants at Lund University, Agrifood Economics Centre and EALE 2017 in S:t Gallen, Switzerland.

1 Introduction

In countries worldwide, governments are committing to large scale investments in high-speed broadband infrastructure (OECD, 2011). Both the European Commission and the U.S. Federal Communications Commission (FCC) have set ambitious targets stating that by 2020, half of all households should have access to at least a 100 Mbit/s connection, a speed which can only be achieved via new information and communications technology (ICT) infrastructure based on optic fiber.

While policymakers often claim large benefits of upgrading networks to give more people access to a fast broadband connection (Kenny and Kenny, 2011), the effects of widespread broadband adoption are poorly understood and require further research. As many countries are currently transitioning to next-generation broadband networks (see e.g., OECD Broadband Portal for cross-country access statistics), understanding the costs and benefits of widespread high-speed broadband adoption is crucial.

This study is part of a growing literature on the socio-economic consequences of broadband. Several empirical studies have employed differences in the timing and location of broadband roll-out to estimate causal effects. To our knowledge, ours is the first empirical study of the effect of high-speed broadband access at home on educational outcomes. We examine the effect of high-speed broadband at home on the academic performance of upper secondary school students in Sweden. We use the differing roll-out of high-speed broadband via optic fiber across localities (parishes), and the resulting within-student variation in broadband exposure, in order to identify the causal effect of broadband on upper secondary school grades. Our register data comprise detailed records on school achievements for all upper secondary school students in Sweden up until 2012. To measure local levels of broadband access, we use an annual nationwide survey conducted among Swedish internet service providers (ISPs) by the Swedish Post and Telecom Authority (PTS), providing high-resolution spatial data on locally available access techniques which we aggregate to the parish level. Our results indicate a small but robust negative effect of high-speed broadband on GPA. Our preferred estimate indicate an average effect size about 5 percent of a standard deviation, with boys, low-ability students and children to parents with low education suffering the largest negative effects.

So far, there have been few studies examining the effect of broadband on educational outcomes. Among the best is the rigorous work by Faber et al. (2015), exploiting the fact that distance to the nearest telephone exchange determines the speed of copper-based DSL (digital subscriber line) broadband. They use a regression discontinuity design (RDD), with residential distance to the nearest exchange as the running variable and the as-good-as-random geographical borders between telephone exchange catchment areas as thresholds. They report a zero effect on standardized test scores, but also show that crossing a boundary seems to have a limited effect on access speed, meaning that the marginal speed increase could be too small to impact behavior. A related area of research is how computer use affects student achievement. Vigdor et al. (2014) analyze the effect of home computer use on American students' standardized test scores in school grades 5-8. Using a student fixed effects model similar to that employed in the present study, they show that a home computer has a small, but significant, negative effect on math and reading scores. Using the number of ISPs connected to a local node as a proxy for broadband coverage, they also show that broadband access reduces homework effort and seems to widen racial and socio-economic gaps.

Using an RDD to examine the effects of a home computer voucher program directed at low-income households in Romania, Malamud and Pop-Eleches (2011) demonstrate that home computer use decreases children's grades, but increases their computer skills. In that study, children reported using the computers not for educational purposes, but rather to play games. However, the share of households with access to the internet in the study was low and the program did not seem to have any effects on internet use. Broadband expansion has also been associated with decreased teen fertility (Guldi and Herbst, 2016), with decreased sexual activity suggested as a mechanism. Aguiar et al. (2017) show that online activities can crowd out offline leisure and labor supply. Using a structural approach, they show that improved leisure technology (e.g., broadband) causes young men to reduce hours worked in favor of gaming and other recreational computer use. This shift can explain 40-80 per cent of the decline in hours worked compared with older men since 2004. Turning to labor demand, there is evidence that ICT complements the skills acquired by a formal education (Akerman et al., 2015). In this paper we study the effects on grades. However, it is important to recognize that effects of ICT may go beyond conventional measures of human capital formation and that certain cognitive skills may benefit from ICT use.

The remainder of the paper is organized as follows. The next section provides a primer on internet access technologies. Section 2 presents a simple conceptual model and provides some evidence on changes in student time-use. Section 3 and describes our data. Section 4 presents our empirical specification and our results. Some concluding remarks are given in section 5.

1.1 Internet access technology

Broadly speaking, there have been three generations of internet access technologies. With the advent of the internet and home computers came the dial-up modem. Using existing telephone lines and dial-up modems, internet access came at a low cost to consumers and ISPs. Applications such as e-mail, online chats, and browsing became common.

The second generation – what became known as 'broadband' – also entered homes over existing infrastructure, specifically phone lines (for DSL) or copper cables used for cable television. With new modems and upgrades to operator node networks, access speeds increased by orders of magnitude compared with dial-up. Access to DSL or cable broadband has now become commonplace in many developed countries (the OECD average is currently around 25 DSL and cable broadband subscriptions per 100 inhabitants). However, the infrastructure underpinning the first- and second-generation access has many limitations, stemming from the fact that the copper cables are not suited to carry high-frequency signals, limiting access speed and reliability. For example, DSL speeds quickly decrease with the distance between the consumer and the operator node. In the latest generation of broadband, copper cable is fully or partly replaced with optic fiber, enabling further increases in access speed. This study focuses on the rollout of broadband delivered via optic fiber directly to consumers' homes, a system also known as 'fiber to the home' (FTTH).

A distinguishing feature of our study compared with previous studies on the socio-economic consequences of broadband is the margin of access speed. Most other studies have focused on the effects of going from a dial-up modem to a DSL or cable connection. By 2007, the first year for which we have data on local availability of broadband, the share of the Swedish population with DSL access is reported at 97.8 percent (Swedish Post and Telecom Authority, 2008)[1]. Thus, our baseline household has access to a conventional copper-based broadband connection and at the margin, the vast majority of households in our study are transitioning from broadband to faster broadband. It is important to note that, in parallel to the expansion of FTTH, the DSL and cable connections are also upgraded by replacement of copper cable with fiber, as the distance that the signal has to travel over old telephone lines or coaxial cables is reduced. Therefore, our measure of fiber coverage also captures upgrades to older technologies, in addition to FTTH.

What are the effects of going from a copper-based connection to fiber? Or, what difference does a super-fast internet connection make when the user already has a fast connection? To quantify the speed increase associated with a fiber upgrade, we use data from a Swedish NGO that provides an online tool for measuring consumer access speeds (The Swedish Internet Foundation, 2013). Using 96 million measurements taken between 2008 and 2013, it puts the average speed of DSL at 11 megabits per second (Mbps), compared with the 57 Mbps of a FTTH connection. It also reports a 50 percent reduction in average latency times when going from DSL to fiber.

How does an increase in access speed affect consumer usage? A simple comparison of internet use between consumers on slow and fast connections will be biased, due to self-selection into different plans and access technologies. Grover et al. (2016) performs a field experiment to estimate the causal effect of access speed on internet use. Working with a large American ISP, those authors randomly upgraded households currently on a 100 Mbps plan to 250 Mbps, without informing the households. Despite already having access to a fast connection, they found that data volumes of these households increased relative to the control group. Interestingly, households that were not fully utilizing the available capacity prior to the upgrade experienced the largest relative increase in demand, suggesting that consumers either started to use more bandwidth-intensive services and/or increased the time spent on internet use.

From a technical perspective, the marginal cost to the provider of supplying additional bandwidth is assumed to be small. Due to better profit margins on high-bandwidth plans, ISPs have an interest in convincing the consumer to upgrade their connection. With this in mind, it is interesting to examine their main selling points. Google (2016) states that fiber means "less time buffering videos [and more] online gaming". Swedish ISPs present similar arguments. Fiber is described as improving online video streaming (often presenting scenarios where multiple family members are watching different video streams at once). The online video-streaming service Netflix recommends a 25 Mbit/s connection

[1]This number reflects households living at an address where at least one ISP reported able to supply a DSL connection "without incurring a significant cost". Actual take-up was reported at around 70 percent in a government survey (Statistics Sweden, 2008).

to stream high-definition content, well above the average Swedish DSL connection. Online gaming and downloading large files (the reference to online piracy is never explicit, for understandable reasons) are also big selling points. The aforementioned activities are all bandwidth-intensive in the sense that they benefit from increases in bandwidth and/or decreases in latency times. As an example, downloading a 3-gigabyte movie takes about 40 minutes using an average DSL connection, while with a 50 Mbps fiber connection the time is reduced to 8 minutes. To some, this reduction may be of little importance. However, a decrease in the time between deciding to watch a movie and pressing play may influence the consumption decision. Another factor relevant for bandwidth constraints is household size. With multiple people sharing available bandwidth, the marginal effect of fiber should increase.

2 Conceptual framework

Given previous research, we argue that faster broadband is more likely to act as a leisure enhancing technology rather than e.g. as a complement to studying. Aguiar et al. (2017) models ICT as a shift in the demand for leisure, causing individuals (males in particular) to reduce labour supply in favor of more free time. As adolescents below the age of 18[2] have limited options outside of formal education and are provided for by parents or legal guardians, their trade-off is better characterized as one between leisure and the future return to human capital rather than one between leisure and current consumption.

In an intertemporal setting, it is intuitive that an increased value of leisure today (due to e.g. ICT) would cause a reduction in labor supply and/or time spent of acquiring human capital, ceteris paribus. While there may be valid paternalistic arguments that adolescents are not able to "rationally" allocate their time between leisure and schoolwork, thus motivating monitoring or other interventions, it is intuitively appealing to model the trade-off as a self-control problem. Thaler and Shefrin (1981) models self-control in a framework where the individual consists of two separate agents, the short-term "doer" acts to maximize instantaneous utility while the long-term "planner" is concerned with lifetime utility.

Building on this concept, Fudenberg and Levine (2006) develops a tractable model of the dual self, incorporating a mechanism by which the planner can control the actions of the short-term self. We adapt their framework to sketch a simple two period model. In the first period (adolescence), the agent has preferences over leisure (h_1). Time not devoted to leisure is spent on acquiring human capital, which is transformed into consumption in the second period (adulthood). To keep things as simple as possible, we abstract from the second period labor supply decision and assume that the individual has preferences over a single consumption good c in the second period. We formulate the problem of the long-run self as

$$\max_{h_1} \quad u_1(\theta h_1, r) + \delta u_2(c(l_1)) \tag{1}$$

$$l_1 = 1 - h_1 \tag{2}$$

$$0 \leq h_1 \leq 1 \tag{3}$$

[2]18 is the age of majority in Sweden

where potential second period consumption c is increasing in the first period investment in human capital, l_1. $\theta > 0$ is a technology shifter in the utility of leisure similar to Aguiar et al. (2017). The first period short-run self only cares about instantaneous utility, and solves the trivial single-period problem

$$\max_{h_1} \quad u_1(\theta h_1, r) \tag{4}$$

$$0 \leq h_1 \leq 1 \tag{5}$$

The actions of the short-run self can only be altered by the long-run planner using a costly self-control action r which fulfills the property $u_1(\theta h_1, r) \geq u_1(\theta h_1', r)$ for all h_1'. This property means that given r, the outcome chosen maximizes short-run utility. Following Fudenberg and Levine (2006), we assume that the cost of self control is proportional to the difference between the unconstrained outcome chosen by the short-run self and the outcome chosen under self-control (evaluated at $r = 0$):

$$C(h_1) = \gamma[\max_{h_1'} u(\theta h_1', 0) - u(\theta h_1, 0)] \tag{6}$$

where γ is a scaling parameter for the cost of self-control. The short-run self maximizes utility by choosing $h_1 = 1$, meaning that $\max_{h_1'} u(\theta h_1', 0) = u(\theta, 0)$. Since the cost is expressed in terms of utility, we can rewrite the long-run problem with costly self-control as

$$\max_{h_1} \quad u_1(\theta h_1, 0) - \gamma[u(\theta, 0) - u_1(\theta h_1, 0)] + \delta u_2(c(1 - h_1)) \tag{7}$$

Assuming that $c(1 - h_1)$ is simply $1 - h_1$ and that the individual has the instantaneous utility function

$$u(\theta h_1, 0) = \frac{\theta h_1^{1-\eta}}{1-\eta} \tag{8}$$

where $\eta > 0$ determines the rate at which marginal utility diminishes. We can solve for h_1 using the first order condition:

$$h_1 = \left((1+\gamma)^{-\frac{1}{\eta}} \theta^{\frac{\eta-1}{\eta}} \delta^{\frac{1}{\eta}} + 1\right)^{-1} \tag{9}$$

As faster internet access increases the quality and number of leisure activities, we can think of high-speed broadband as both increasing the cost of self-control (an increase in γ) and well as shifting demand for leisure upward (an increase in θ). As shown by (9), h_1 is increasing in γ, meaning that as self-control becomes more costly, the long-run self will find that it is optimal to "give in" to the short-run self and allocate more time to leisure. This prediction is consistent with the findings of Aguiar et al. (2017). Faber et al. (2015) makes a similiar argument and refers to the substitution from productive work to leisure due to ICT as "the Facebook effect". The effect of θ depends on how rapidly the marginal utility of leisure diminishes.

The appeal of the dual-self framework lies in capturing the sometimes ambiguous attitude towards ICT as a leisure enhancing technology. While the lifetime utility of the long-run self decreases as the cost of self-control goes up, the short-run self enjoys greater utility of leisure. While there is little doubt that people enjoy the benefits of faster internet, there is emerging evidence that ICT can reduce well-being (Twenge et al., 2018) and that people may try to commit to spending less time online by e.g. exiting social media (Schoenebeck, 2014). The latter can be viewed as a self-control action. Rather that relying on the short-term self to simply log in less frequently, one can preemptively restrict future outcomes by deactivating or deleting one's account.

Parental supervision can be an alternative to self-control. We use parental education as a proxy for the parential resources [3] when examining heterogeneity in the effect of fiber. However, it is important to recognize self-control and patience as inheritable traits, meaning we cannot truly disentangle the role of parental supervision from confounding characteristics.

We also consider non-linearities and spillovers in the effect of broadband expansion. As a consumption goods, high-speed internet has potential network externalities, meaning that it's appeal is increasing in the number of my peers who use it. This may produce a non-linear effect where the marginal effect of additional coverage is dependent of the level of initial coverage. We test for a non-linear effect by conditioning the the initial level of coverage.

Spillovers may bias our estimates if students from areas with low coverage can take advantage of broadband in high coverage areas via peers. Since the result is that non-treated students are partially treated, cross-parish spillovers would bias our estimates towards zero. On the other hand, within parish spillovers can exacerbate the effect of expanding coverage if take-up affects peers who's households are not yet covered. We consider this spillover part of the "true" effect, but this also speaks to the idea of a non-linear effect.

2.1 Trends in student time use

Our simple model predicts that if faster internet access increases the quality of leisure, we should expect a shift in time-use towards leisure.[4] This shift away from e.g. time spent on school work is perhaps the most obvious mechanism by which broadband could affect academic achievement. To examine the potential scope of this mechanism, we turn to data on student time-use. As Internet usage patterns are rapidly evolving, we would ideally want to limit ourselves to time-use data dating from our main study period which is 2007–2011, but as data availability is limited we have to extend this window a few years.

An annual survey of online habits in Sweden puts the share of daily internet users among 16- to 24-year-olds at above 90 percent (IIS, 2011). The same survey documents a distinct gender differential. For example, 28 (9) percent of boys (girls) ages 16-25 report daily online gaming. Another survey around the same time corroborates the gender differences in usage patterns ((Swedish Media Council, 2013)); 50 percent of responding 15 year-old boys spent at least 3 hours per day on computer games. The corresponding number for girls was 0 percent.

In 2011, researchers working with the file-sharing site 'The Pirate Bay' conducted a survey among visitors to the site (Svensson et al., 2013). Of the 2000 respondents based in Northern Europe aged 17 or

[3] Education is known to be correlated with several socio-economic characteristics of the family. See Björklund and Salvanes (2011) for a review of the literature on education and family background.

[4] Vigdor et al. (2014) examines time-use and the link to internet at home. They interpret their findings as time-spent on homework becoming less efficient as a result of the distraction from broadband.

Table 1: Trends in student time use, PISA surveys

	Internet use, weekday (h)	Internet use, weekend (h)	Internet use, at school (h)
2015 (second wave)	1.186***	1.351***	1.092***
	(0.0603)	(0.0635)	(0.0507)
Boy	0.812***	1.067***	0.0885*
	(0.0630)	(0.0664)	(0.0530)
2015 × Boy	-0.738***	-0.781***	-0.287***
	(0.0862)	(0.0908)	(0.0724)
Outcome mean	3.58	4.23	1.48
Observations	9,202	9,198	9,203

The outcome in columns 1 and 2 is the answer to the question *"During a typical weekday (weekend), for how long do you use the Internet outside of school?"*. In column 3, the question is *"During a typical weekday, for how long do you use the Internet at school?"*. We linearize the categorical response to hours by imputing the mid-point for each category. Standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

younger, about 4 percent were girls. Among daily file sharers, girls made up just over 2 percent of the sample. However, adolescent girls are avid users of other online services, e.g., they are over-represented when it comes to music streaming and social media use (Swedish Media Council, 2015). However, these activities do not require a lot of bandwidth and consequently do not benefit from faster access speed to the same extent as the services where boys make up the majority of users. The fact that bandwidth intensive Internet use seems more prevalent among boys leads us to hypothesize that any negative effect of broadband on academic achievement should be greater in magnitude for boys than for girls.

As the roll-out of high-speed broadband progresses, do adolescents spend more hours online? The Program for International Student Assessment (PISA) is an international triennial survey of the skills and knowledge of 15-year-old students within reading, mathematics and science. The latest waves included an ICT questionnaire where the student is asked about attitudes towards ICT, usage and familiarity. Each wave includes a representative sample of about 5000 Swedish 15-year-olds, most attending the final year of lower secondary school (9th grade), i.e. the year before they start upper-secondary school. The PISA survey unfortunately does not report student or school location, so we cannot make inferences relating to local broadband expansion.

Table 1 reports some descriptive regressions from the PISA survey in 2012 and 2015. In these two waves, the questionnaire explicitly asked for total time spent online during weekdays, weekends and in school. We linearize the categorical response by imputing the mid-point for each category.[5] The estimates are reported in columns 1-3. Daily internet use among respondents increased by about an

[5] E.g. the response category "0-30 minutes" is coded as 15 minutes.

Figure 1: Trends in online activities, PISA 2006-2015

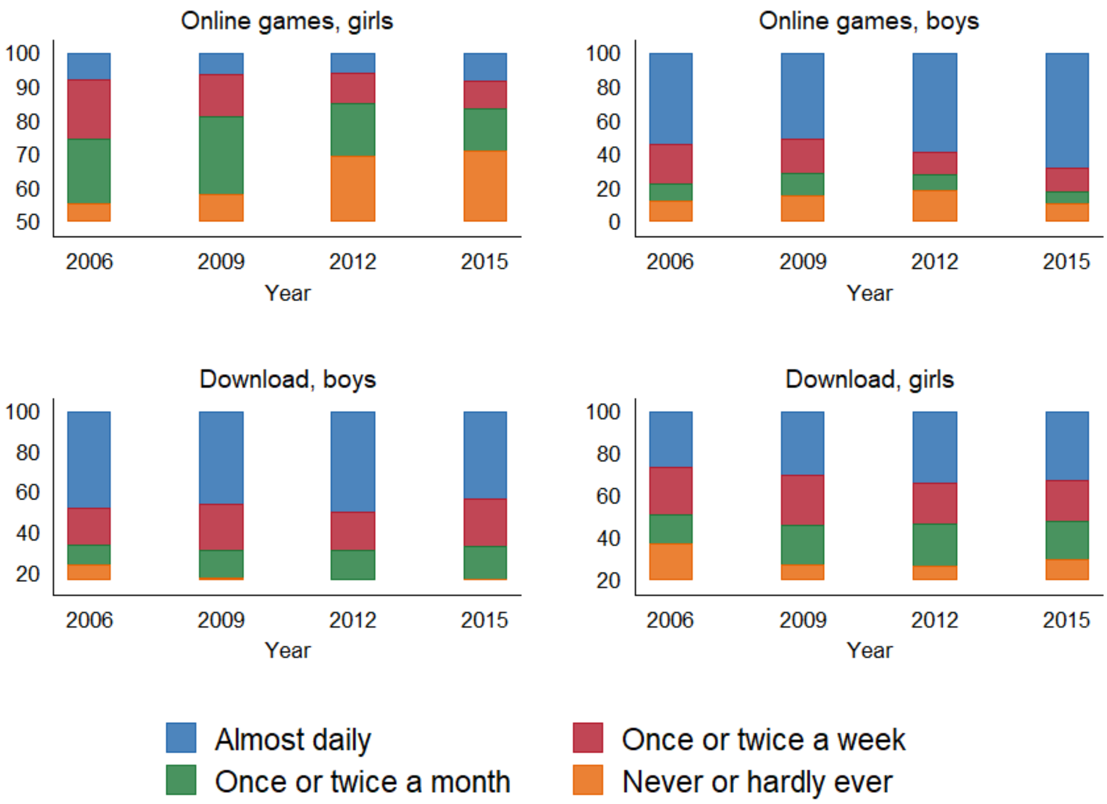

hour between 2012 and 2015, both in and outside of school. There was a clear gender differential in 2012, but the interaction between 2015 and gender reveals that the gap in time-use is closing, and is even reversed for internet use at school. Since 2006, the questionnaire asks how frequently the student uses the Internet for specific tasks, including online games and downloading games/music/movies. Figure 1 details trends in the categorical responses to these questions. For girls, we note a slight upward trend in using the internet to download but a substantial decrease in gaming. For boys, downloading seems flat and gaming seems to be gaining in popularity. The lack of clear shifts between 2012 and 2015 suggests that the usage increase presented by table 1 could be driven by other online activities such as social media or video streaming. Changes in activity mix could also explain the closing gender gap.

Since time-use is a zero sum game, what are students spending less time on as Internet use increases? An OECD report comparing time spent on homework as reported in the 2003 and 2012 PISA surveys noted a small (about 20 minutes per week) decrease for Swedish students (OECD, 2014).[6] Swedish students were consistently in the bottom end of the distribution in terms of total time spent on homework.

3 Data

Since 2007, PTS conducts an annual survey of telecom operators and ISPs regarding broadband access in Sweden (Swedish Post and Telecom Authority, 2008). This information is matched to register data on location, transcripts from upper secondary school, and other background information. The respondents in the PTS survey are asked to produce a list of all the addresses where they supply a connection to consumers. As responding to the survey is considered mandatory, more than 90 percent of telecom operators complete the survey every year. A single grid square is considered to be covered if at least one building within the square has access to fiber, i.e., each 250 m by 250 m grid is either covered or not. Note that a consumer located in a covered square could face a significant cost of actual take-up should they live far away from the ISP node.

Using data on the working-age (16-64) population within each 250 m by 250 m square[7], we calculate a measure of coverage at parish level by weighting the coverage in each square by its population (see Figures 2 and 3 for a visual representation). Our measure of fiber coverage can thus be interpreted as the share of the parish population covered or, equivalently, as the probability of being covered conditional on the parish of residence. A small share of students (about 1 percent of the sample) live in parishes where the recorded change in coverage between 2007 and 2011 is negative. In one case, we confirmed that this is due to a reporting error by the ISP. Consequently, we excluded these students from our sample. The 1387 parishes in our data are on average home to about 7000 people and cover an area of about 300 km2.

[6] Questions relating to time spent on studying outside of school have been reframed between PISA waves, making it very difficult to make intertemporal inferences between e.g. 2012 and 2015. Differences between waves include reporting time spent by subject or in total, time spent alone or with a tutor and continuous or categorical responses.

[7] We use 2013 data on grid population, as this is the earliest year available to us.

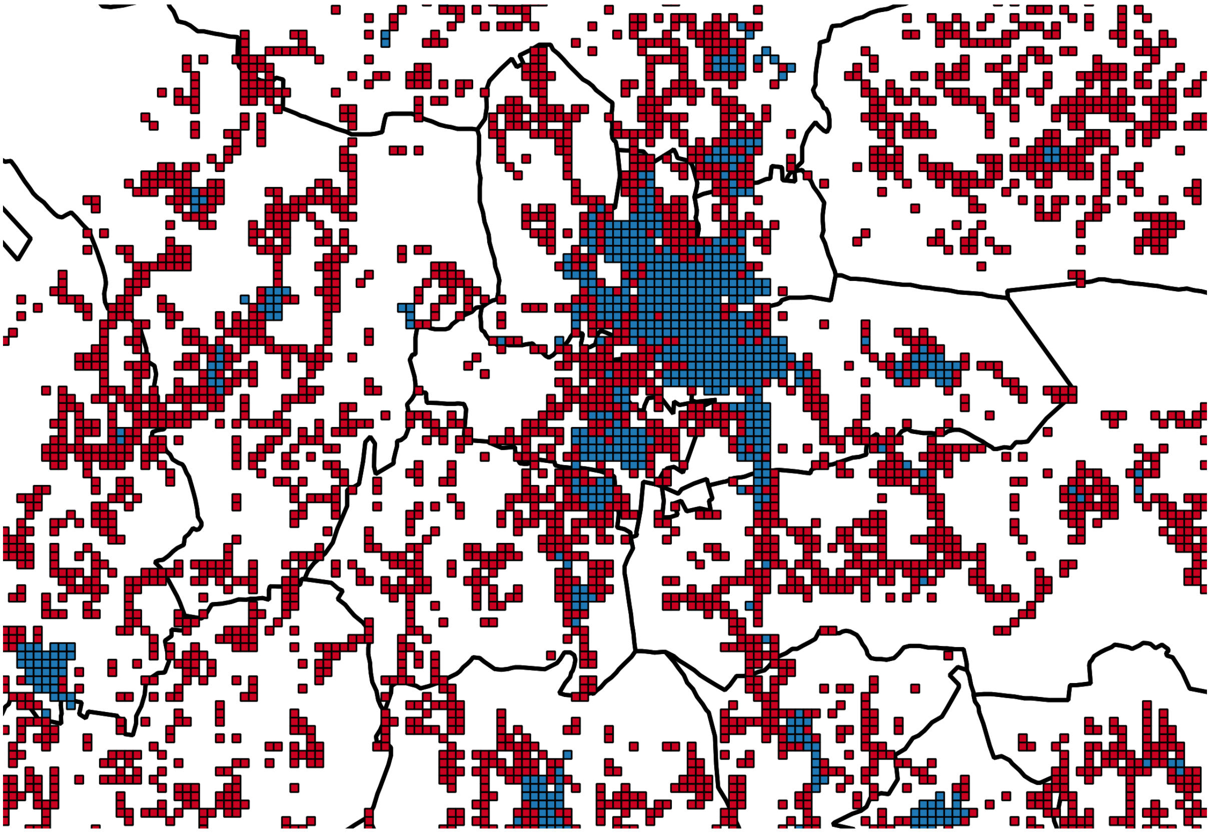

Figure 2: Raw 250*250 meter grid data. Blue areas are covered, white areas are unpopulated.

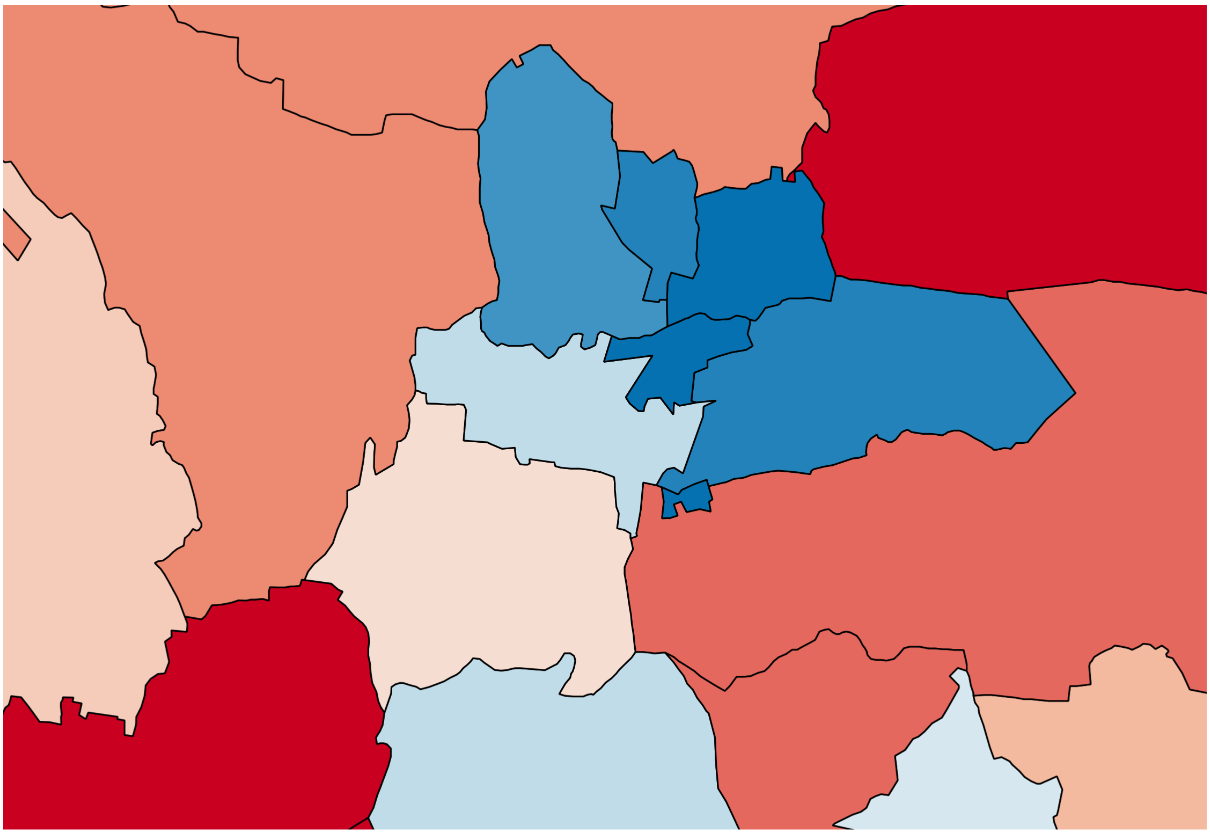

Figure 3: Parish measure of fiber access. Areas in blue have higher coverage.

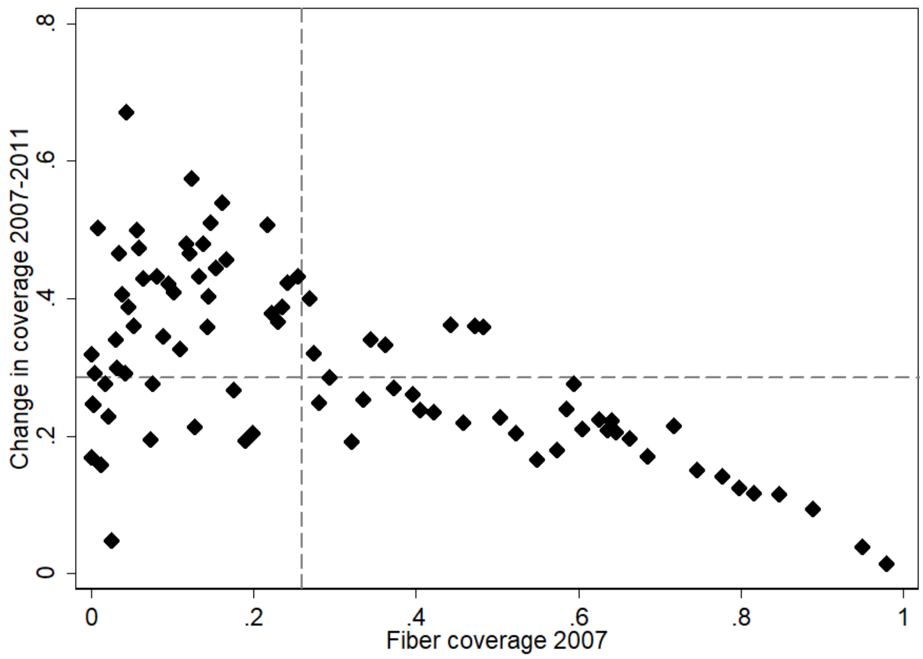

Figure 4: Change in fiber coverage plotted against the initial level of coverage, averages within 87 equal-size bins. Dashed lines represent sample means.

Figure 4 shows the variation in our sample in terms of treatment intensity (change in fiber coverage between 2007 and 2011) and the initial level of coverage (in 2007). To reduce visual clutter, we bin our data into equal-size bins and plot within-bin means. As can be seen, almost everyone in our sample is treated to some extent, but there is variation in treatment intensity even for a given level of initial coverage.

From register data, we have detailed transcripts on all students graduating upper secondary education up until 2012 . Our data include GPA at graduation and grades for all individual courses. An important consideration for our study is that students in Swedish upper secondary school are graded continuously throughout their three years of study. Subjects are typically taught over multiple courses, and courses taken during the first year carry the same weight as courses taken during the last year when calculating GPA at graduation. GPA is calculated as a course credit-weighted average of grades and presented as a number from 0 to 20, where 10 is the equivalent of obtaining a pass grade in all courses taken. While we do not know when a course was completed, course codes (e.g., MA1201 for introductory math and MA1204 for calculus) provide some information regarding when a course was taken. All students in upper secondary school take a number of mandatory introductory courses ('core courses') equivalent to about a full year of studies. During the years in our study, these courses consisted of introductory Swedish, Mathematics, English, Arts & Crafts, Physical Education, Religion, Science, and Social Studies.

While the exact timing of the courses is not regulated, we rely on the fact that completing the introductory course is formally recommended before taking more advanced courses in the same subject, meaning that the curricula will typically have students completing the core courses early on. Moreover, if the student passes the course, the grade cannot be revised. Vocational programs typically give practical courses within the chosen field as well as provide provide apprenticeships for their students during the latter two years of secondary school, meaning that they are also likely to complete the core part of the program early on. Consequently, we use GPA on core courses as a proxy for first-year GPA, and the GPA based on non-core courses as a proxy for second- and third-year GPA. Curricula from school websites available at the time of writing confirms that these are accurate proxies.[8] Within a student, the measurement error caused by proxying the timing of grades has must have mean zero (a positive measurement error in first year GPA will be canceled out by a negative error in latter year GPA). While we could be systematically mis-measuring GPA within a period, this would only be a problem if student curricula is somehow correlated with local fiber roll out, which seems unlikely given that students within a school typically reside in parishes with varying degress of broadband expansion.

To identify a causal effect, we exploit within-student variation in fiber coverage between the first and second/third year. We associate first-year GPA with coverage in fall of the first year of upper secondary school, and later-year GPA with the average coverage in fall of the second and third year. For our analysis, we standardize GPA to have zero mean and unit standard deviation.

To exploit the large cross-sectional differences in fiber roll out, we require data on location. We use data from the Swedish pharmaceutical register and the national patient register, which records the parish of residence each time a prescription is processed by a pharmacy and during outpatient visits (not including primary care visits). An interesting feature of the pharmaceutical register is that in many cases it reports a location multiple times during each year, allowing us to minimize measurement error due to students moving e.g., to take up university studies in fall in graduation year. While this register only provides information on students who have been in contact with health care providers, through a sequence of matches based on the prescriptions and health care visits by the individual, younger siblings, and parents (see table A.2), we can link 96 percent of students to a parish during their upper secondary school years.

However, attrition is unlikely to be independent of student characteristics and we note that the sample without location differs significantly from the main sample with regard to several observables (see table A.1).[9] As shown in table 2, parish coverage increases by about 50 percent (15 percentage points) between the first and last two years. GPA also increases between year 1 and the later years, which could be due to several reasons. It could reflect the effect of initially having to take 'core courses' that the student might not find very interesting and only taking courses within their chosen track later, but might also be the result of students maturing or the formation of student-teacher bonds.

Turning to GPA and gender, the gap between girls and boys increased by about 4 percent between 2003 and 2012, as Figure 5 shows. Note that in order to better illustrate the change in the gender gap, mean GPA for both boys and girls is normalized to zero in 2003 in Figure 5. The absolute gender gap

[8]Unfortunately, we have not been able to locate any archived curricula or schedules from our study period.

[9]Boys are over-represented in the group without location, likely because Swedish girls are issued medical prescriptions more frequently, i.e. contraceptives.

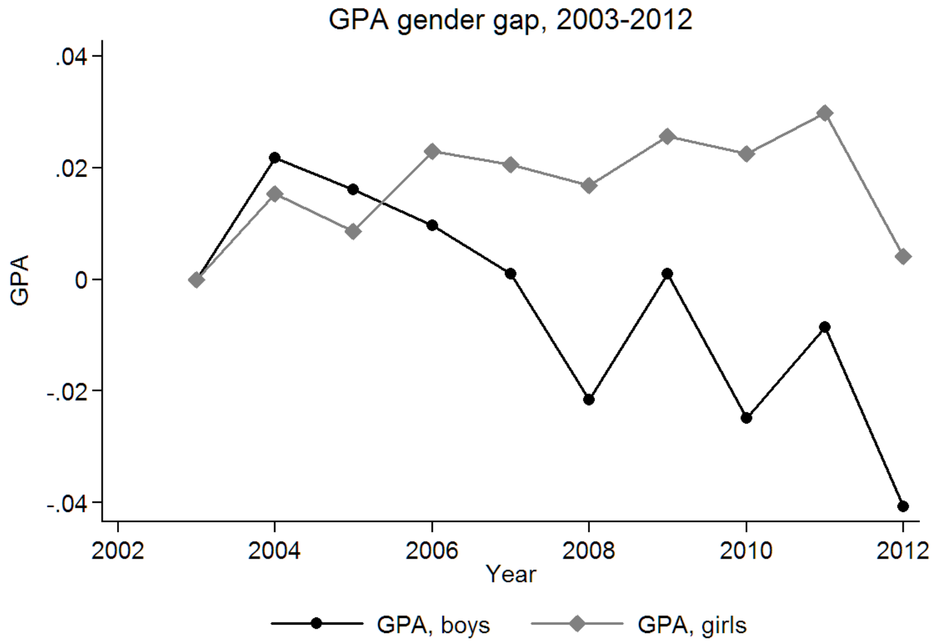

Figure 5: GPA gender gap. GPA is normalized to zero in 2003 for both boys and girls.

during the period is on average 40 percent of a standard deviation.

Our main regression sample consists of students graduating upper secondary school between 2010 and 2012, excluding students who move between their first and third year of upper secondary school (about 10 percent of students). While dropping these students could bias our estimate if the decision to move is related to broadband expansion, our main concern is that we would risk capturing other effects of the move, such as the effects of changing schools, peers, and neighborhood. Another concern is measurement error in our parish variable, e.g., if a student's parents are separated, we might match their location based on the father's residence in one year and the mother's residence in the next. It is only when we consistently match a student to the same parish during all three years of upper secondary school that we can be reasonably sure that we are observing the true parish of residence. However, all our results are robust to including movers. Since we require data on fiber coverage when starting upper secondary school, our earliest cohort consists of students starting in fall 2007 (i.e., graduating in summer 2010).

4 Results

For our individual fixed effects specification, we assume that the relationship between GPA and fiber coverage for individual i living in parish p in year t can be described as

$$GPA_{ipt} = \beta Fiber_{pt} + \alpha_i + \gamma_{pt} + \varepsilon_{ipt} \tag{10}$$

For $\hat{\beta}$ to represent an estimate of the causal effect of fiber on GPA, we require that, conditional on time

Table 2: Summary statistics by change in fiber coverage

	Full sample	Bottom 50 % of fiber rollout	Top 50 % of fiber rollout
Year of birth	1991.8	1991.8	1991.8***
Boy	0.49	0.49	0.49
GPA at graduation	14.1	14.0	14.1***
GPA, year 1	13.8	13.5	13.8***
GPA, years 2-3	14.1	14.1	14.1
GPA, 9th grade	223.0	220.7	223.7***
Fiber coverage, year 1	0.30	0.22	0.33***
Fiber coverage, years 2-3	0.42	0.23	0.47***
Academic program	0.47	0.43	0.48***
Immigrant	0.05	0.03	0.06***
Mother's years of schooling	11.9	11.7	11.9***
Father's years of schooling	11.8	11.5	11.9***
Mother's income (log SEK)	11.5	11.4	11.5***
Father's income (log SEK)	11.9	11.8	11.9***
Number of students	250,433	56,827	193,606 .

***: Difference with respect to the second column is statistically significant at the 1 percent level.

and individual fixed effects, students in less treated parishes represent an unbiased estimate of the counterfactual GPA difference for more treated students (the parallel trends assumption). Formally:

$$E\left[GPA_{0ipt}|\alpha_i, t, Fiber_{pt}\right] = E\left[GPA_{0ipt}|\alpha_i, t\right] \tag{11}$$

where GPA_{0ipt} denotes the (counter-factual) GPA of individual i in the absence of fiber expansion. It can be useful to think of our specification as the reduced form of an instrumental variables specification, where coverage at parish level is used an instrument for actual household take-up of fiber (which we do not observe). While the decision to connect one's home to the local fiber network is endogenous, short-term changes in parish coverage are arguably exogenous from the student's perspective, and a valid instrument for household take-up of fiber. Our estimates of β represent intention-to-treat (ITT) estimates, meaning we estimate the effect of gaining access to fiber.

Our difference in differences design is more complex than a textbook implementation for two reasons. First, treatment is continuous. Different parishes experienced different expansions of coverage during our window of study. Second, our coverage data does not cover the roll-out that occurred prior to 2007. The mean parish coverage across all students in 2007 was around 25 percent (see figure 4). Unconditionally, fiber roll-out is not randomly assigned. In Table 2, we document differences between the 'treatment' and 'control' groups for a number of observables. The differences in parental education and income suggest that students who receive more intensive treatment are positively selected. As a balancing test, we aggregate data to parish means and test if pre-determined parental cohort, income and education can predict student fiber coverage when we include area fixed effects (see table 3). Once we condition on parish fixed effects and trends as well as municipal by cohort fixed effects (column 4), the explanatory power of parental observables disappears.

In order to visualize the pre- and post-treatment GPA trends, we sort students into 'control' and 'treatment' groups by splitting the parish-level distribution of the change in fiber coverage between 2007 and 2011 at the median, i.e. splitting the sample by treatment intensity with those in the top half receiving greater coverage expansion. The main threat to identification is that the within student difference between first- and $2^{nd}/3^{rd}$ year GPA is affected by some omitted variable correlated with fiber expansion. Because each student attends upper-secondary school only once, we cannot assess pre-trends in outcomes at the student level. However, we can compare the parish mean of the within student GPA difference. We calculate the within student difference between first- and $2^{nd}/3^{rd}$ year GPA by graduation cohort and take means over over 'control' and 'treatment' groups for each cohort. As figure 6 shows, the trends are largely parallel and do not begin to diverge until we get to the cohort graduating in 2010. By construction, this is the first 'treated' cohort in our sample. Although cohorts prior to 2010 experienced broadband roll-out, this does not seem to have put them on diverging trends. This indicates that the less treated group represents a valid counterfactual. The clear difference in levels could be the result of a different track mix between the two groups[10]. When we perform the same exercise but split parishes by the *initial* (2007) level of coverage, we find a slight indication of an earlier divergence in trends, possibly

[10]There is a constant 5 percentage point difference in the share of students choosing an academic track (see figure A.2)

Table 3: Balancing test, predicting fiber by parental characteristics

	(1)	(2)	(3)	(4)
Mother's schooling	0.0176	-0.00603*	-0.00550*	-0.00174
	(0.0123)	(0.00355)	(0.00331)	(0.00458)
Father's schooling	0.0896***	-0.00734**	-0.00677**	-0.00109
	(0.0122)	(0.00304)	(0.00287)	(0.00412)
Mother's income	-0.0183	-0.00567	-0.0141	-0.00465
	(0.0237)	(0.00846)	(0.00884)	(0.0113)
Father's income	-0.00960	0.00300	0.00121	-0.000699
	(0.0224)	(0.00774)	(0.00808)	(0.00923)
Mother's birth year	0.00697*	-0.00258	-0.00218	-0.000971
	(0.00419)	(0.00175)	(0.00172)	(0.00229)
Father's birth year	-0.00526	0.00286*	0.00165	0.00163
	(0.00376)	(0.00147)	(0.00146)	(0.00208)
F-stat (p-value)	0.000	0.009	0.005	0.927
Cohort FE	YES	YES	NO	NO
Muni FE	YES	NO	NO	NO
Parish FE	NO	YES	YES	YES
Muni $\times$ cohort FE	NO	NO	YES	YES
Parish linear trend	NO	NO	NO	YES
Observations	3,832	3,832	3,832	3,832

The dependent variable is parish fiber coverage. Standard errors clustered at the parish level. *** p<0.01, ** p<0.05, * p<0.1

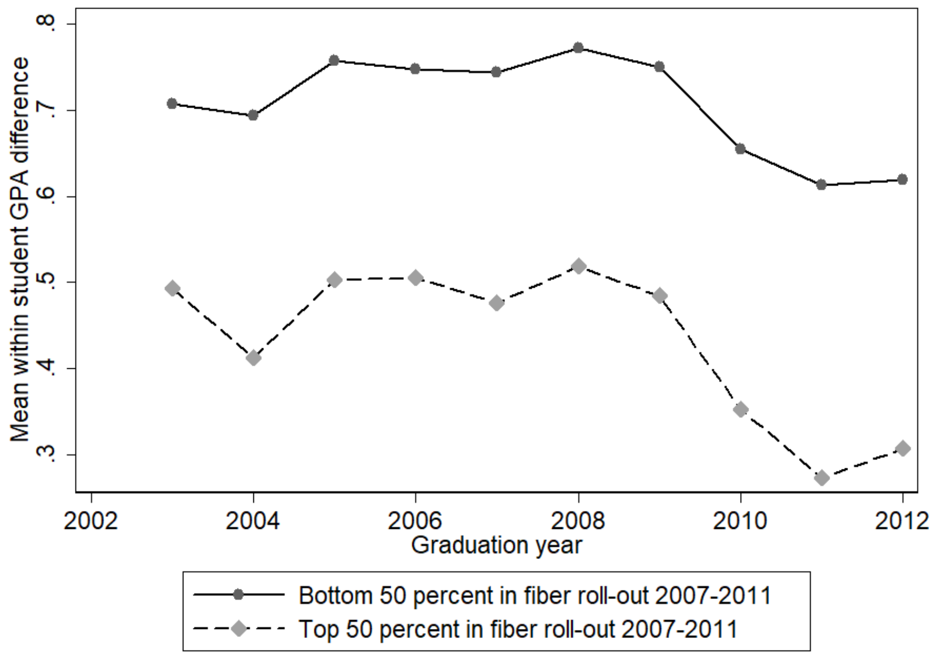

Figure 6: Mean within student GPA difference by treatment intensity.

indicating a treatment effect among the early fiber adopters (see figure A.1).

Table 4 presents our baseline estimates. The first column reports the effects of fiber on GPA at graduation using a rich set of controls. Using within student variation in GPA (column 2) puts the negative effect of a one unit increase in fiber coverage[11] at almost 30 percent of a standard deviation. Adding a linear trend at the parish level drastically reduces the estimated effect (column 3). This is likely explained by the fact that the gradual fiber roll-out has a strong trend trend component. A linear trend explains on average 70 percent of the variation in fiber expansion. By controlling for a linear trend at the parish level, we are eliminating some of the true effect. However, as the pre-2007 expansion may locally jeopardize the parallel trends assumption, we believe that accounting for a pre-existing trend is important. With this in mind, the estimate in column 3 is likely conservative.

Column 4 adds a linear trend at the school level. The significant effect indicates that the results cannot be explained by students in different parishes attending schools with different grading practices or an effect of broadband at school. Our preferred specification in column 5 includes a linear trend at the parish level along with a full set of municipality by year fixed effects. The municipality is an important administrative boundary as they govern not only schools but they are also responsible for administrating national broadband subsidies. This is the specification corresponding to column 4 in table 3. The estimate suggests that complete local fiber expansion decreases student GPA by about 4.5 percent of a standard deviation.

[11]Corresponding to an expansion from from 0 to 100 percent parish coverage

Table 4: Main results

Fiber	-0.0613*	-0.290***	-0.0390***	-0.134***	-0.0445***
	(0.0331)	(0.0299)	(0.0133)	(0.0153)	(0.0146)
Observations	215,975	503,776	503,776	503,776	503,766
Students	215,975	251,888	251,888	251,888	251,883
R-squared	0.449	0.898	0.901	0.909	0.903
Student FE	NO	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	NO
Muni × Year FE	NO	NO	NO	NO	YES
Linear trend	None	None	Parish	School	Parish

Column 1 includes controls for parents' age at birth, education, and income, and the student's GPA in 9th grade and indicators for sex, sibling order, and parish. Fiber is measured as average coverage during upper-secondary school. Columns 2-5 uses within student GPA variation for students graduating 2010-2012. Fiber is measured as coverage during the first year and mean coverage during the 2nd and 3rd year. Standard errors clustered at the parish level. Our sample excludes movers, dropouts, and students in parishes with a negative reported change in fiber coverage during 2007-2011. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

The average treatment effect would be larger than our ITT estimates if actual take-up of high-speed broadband is anything less than 100 percent. Since we do not have a first stage, we can only do a rough back-of-the-envelope estimate of the LATE. A 2012 survey by Statistics Sweden (Statistics Sweden, 2013) puts the share of respondents ages 16 to 24 with a fiber or cable connection at 29 percent. In our sample, mean coverage in 2012 is 55 percent. This gives us a quasi-first stage estimate of $0.29/0.55 = 0.52$ and a LATE of $0.045/0.52 = 0.087$ (using the estimate in column 5).

To put the effect size into perspective, the seminal meta study by Hattie (2015) puts the effect size at a magnitude comparable to sleep deprivation or having your family on government aid. We can also undo the normalization and express the effect in terms of actual grades. A negative effect of 8.7 percent of a standard deviation would be comparable to going from an A to a B in introductory Swedish, English and Math. While this is likely inconsequential to the majority of students, for those eyeing the most selective university programs the impact in terms of future labor market outcomes may be severe.

4.1 Heterogeneity

Table 5 documents heterogeneity in the effect of fiber by way of sample splits. Our data seem to support our hypothesis that boys are affected to a larger degree than girls. As the estimates in columns 1 and 2 show, the point estimate for boys is about 60 percent larger than for girls, with the effect size estimated at 5.3 and 3.3 percent, respectively. As mentioned previously, we observe a 4 percent increase in the GPA gender gap between 2003 and 2012. In 2012, the average parish coverage across all students in our sample was 55 percent. Assuming that fiber had no impact on GPA back in 2003, as Sweden was then still in the very early stages of extending coverage, we can do a back-of-envelope calculation of the impact of fiber on the gender gap. Taking the difference in point estimates and multiplying by 0.55, the differential effect of fiber explains just over 25 percent of the increase in the GPA gender gap between 2003 and 2012.

Table 5: Heterogeneity

Fiber	-0.0526**	-0.0332	-0.0541**	-0.0382*	-0.0213	-0.0465**
	(0.0242)	(0.0208)	(0.0223)	(0.0195)	(0.0205)	(0.0217)
Observations	248,142	255,594	261,798	241,938	236,688	267,058
Students	124,071	127,797	130,899	120,969	118,344	133,529
R-squared	0.893	0.908	0.887	0.914	0.916	0.896
Sample	Boys	Girls	Mother low edu	Mother high edu	Academic program	Vocational program

This table presents the individual FE estimates of equation 1. Columns 1-2 splits the sample by the sex of the student. Columns 3-4 splits the sample by mother's education (< and > 12 years of schooling, respectively). Columns 5-6 splits the sample by the type of program attended by the student. All regressions include municipal by year fixed effects and parish specific linear trends. Standard errors clustered at the parish level. Our sample excludes movers, dropouts, and students in parishes with a negative reported change in fiber coverage between 2007 and 2011. *** p<0.01, ** p<0.05, * p<0.1

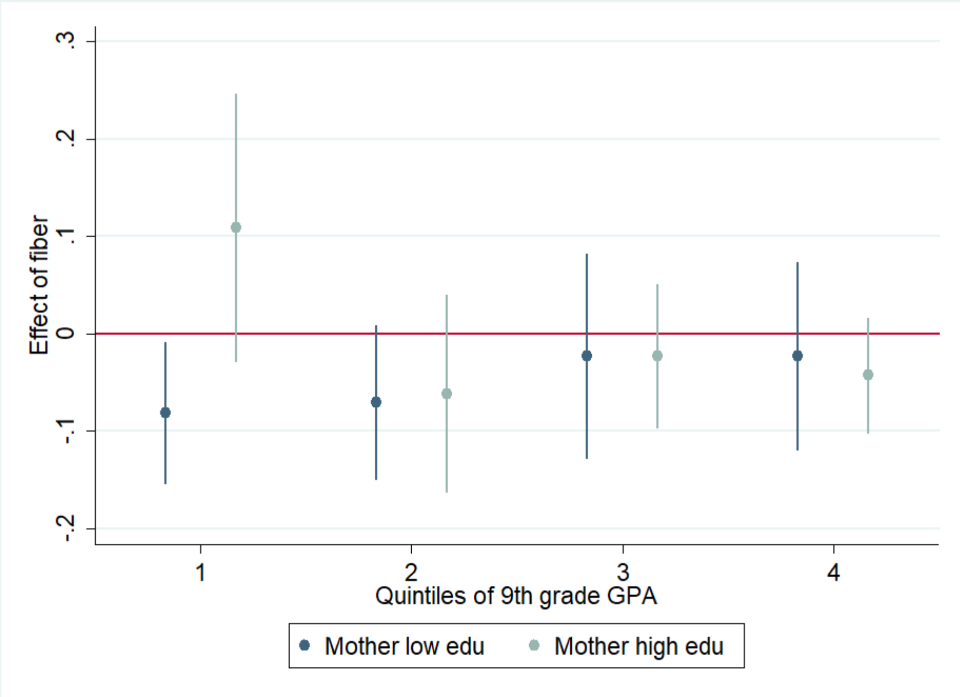

Figure 7: Heterogeneity in the effect of fiber by student ability by mothers' education.

We also explore heterogeneity by parental education (columns 3 and 4 of table 5). Splitting the sample between mothers with more or less than 12 years of schooling [12], we find that the negative effect is concentrated among mothers with low education. This is consistent with parental investments playing a role in mitigating the negative effect, but could also reflect a correlation with student ability and self-control. We find similar discrepancies when conditioning on the father's years of school, but smaller in magnitude (not presented) We also split the sample by the type of program chosen by the student (columns 5 and 6). If students sort into programs by ability, the fact that the effect seems mainly driven by students in vocational programs suggests that low-ability students are hit the hardest.

To disentangle the roles of parental investment and inherited student characteristics, we use GPA in the 9th grade (the final year of elementary school) as a proxy for student ability. To address some of the concerns associated with conditioning on a potential outcome, we split the sample into 9th grade GPA quartiles within each parish to have a prior distribution that is conditional on local broadband coverage and other local characteristics. When we run regressions on GPA quartiles by maternal education (see Figure 7), we find that maternal education is important even when conditioning on prior GPA, but only for students in the lowest quartile.

Finally, we test for non-linearities in the effect by splitting the sample by the initial (2007) level of coverage (see table 6). Column 1 (2) reports the estimate for students in parishes in the bottom (top) of the 2007 coverage distribution. The difference between the two groups is a about 30 percentage points. While we do not find strong support in favor of a non-linear treatment effect, the higher estimate in

[12]12 years of schooling is equivalent to at least a degree from a three-year academic track in upper secondary school. We drop the 10 percent of students missing data on maternal education.

Table 6: Non-linear treatment effect

Fiber	-0.0711*	-0.0444***
	(0.0362)	(0.0163)
Observations	105,754	397,946
Students	52,877	198,973
R-squared	0.896	0.905
Mean initial coverage	0.002	0.32
Student FE	YES	YES
Muni x Year FE	YES	YES

Column 1 (2) conditions on being in the bottom (top) 50 percent of the parish fiber coverage distribution in 2007. Standard errors clustered at the parish level. Our sample excludes movers, dropouts, and students in parishes with a negative reported change in fiber coverage between 2007 and 2011. *** p<0.01, ** p<0.05, * p<0.1

column 1 could possibly be explained by spillovers or a lack of other amenities in rural and less covered areas. The latter could amplify the impact of broadband on leisure quality.

4.2 Robustness

To test the robustness of the results, we first counterfactually introduce fiber 1 to 3 years earlier. The results of this exercise are presented in Figure 8. The counterfactual roll out does not yield any significant effects. However, this does not constitute a clean placebo test, since the roll out we exploit in this paper is likely to be related to earlier broadband investments. Due to lack of data on broadband expansion pre-2007, we can only provide limited evidence of this. In 2000, prior to the first wave of subsidies to broadband expansion, the Swedish government sought to identify rural areas where commercial investment was unlikely to happen (The Broadband Committee, SOU 2000:111, 2000). The earmarked funds were distributed among municipalities in proportion to the estimated cost of extending coverage. Given this subsidy schedule, a larger absolute gap between projected cost and awarded funds should, ceteris paribus, be negatively correlated with broadband coverage. Interpreting the cost-subsidy gap as a proxy for the roll out speed of first-generation broadband, we regress the increase in municipal fiber coverage during 2007-2011 on the cost-subsidy gap (not reported) and find that an additional 10 million SEK in subsidies back in 2000 is associated with an increase in the expansion of coverage during 2007-2011 of 2 percentage points[13], suggesting that the effects of these early subsidies still affected expansion rates and that roll out rates are positively autocorrelated. In addition, using a 2003 PTS survey detailing the number of ISPs that owned local internet infrastructure in Swedish towns and cities, we find that additional

[13]Significant at the 5 percent level

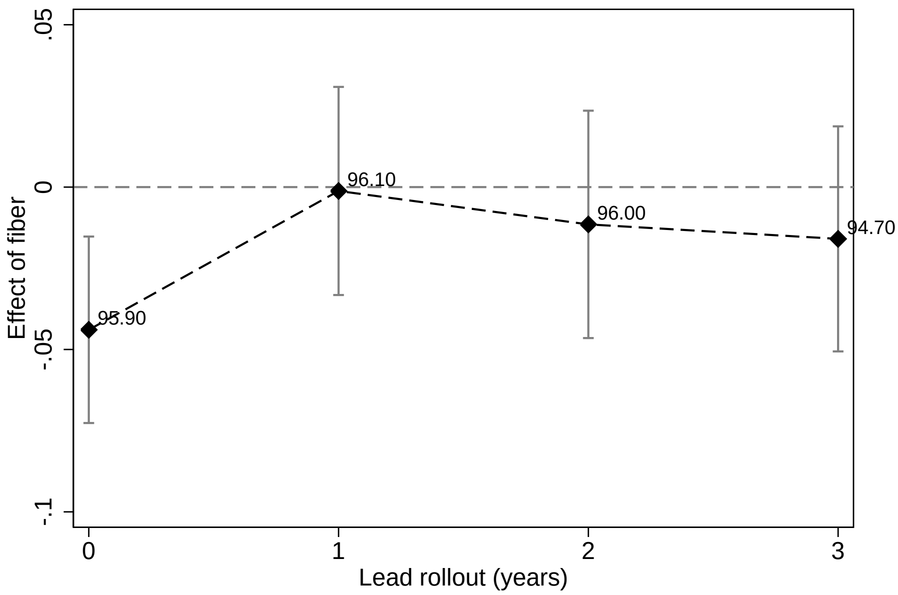

Figure 8: Effect of fiber when the roll out is pushed forward by up to 3 years (e.g., 2004-2008 instead of 2007-2011). The labels denote the share of students that we can match to a parish for all three years of upper secondary school. The "0" estimate corresponds to column 2 of Table 4.

operators in 2003 are positively correlated with broadband expansion during 2007-2011[14].

Nonetheless, with a linear trend at the parish level we are effectively identifying an effect using discontinuous jumps in coverage between two years. These are arguably harder to predict using data on prior expansions and subsidies. Thus, it is reassuring to see that when we counterfactually expand coverage earlier, we do not find any significant effects.

4.3 Effects on dropping out and tertiary eligibility

Next, we examine the effect of fiber along two other margins, dropping out and being eligible for tertiary education [15]. For students who drop out, our records only indicate that they were enrolled but did not receive a diploma at the time of graduate. We do not know any course grades. With only a single data point per student, we resort to a sibling fixed-effects specification where we compare dropout rates between siblings:

$$D_{hpt} = \beta_1 Fiber_{pt} + \beta_2 \mathbf{X} + \alpha_h + \gamma_t + \varepsilon_{ipt} \tag{12}$$

[14]The number of ISPs operating in an area has been used as a proxy for local coverage in previous research (Vigdor et al., 2014; Kolko, 2012).

[15]Students are eligible to apply to universities and university-colleges if they obtain passing grades on at least 90 percent of their course credits, including introductory Swedish, English, and Math.

Table 7: Extensive margins

	GPA	Dropout	Eligibility
Fiber	-0.0286	0.000780	0.0206
	(0.0441)	(0.0140)	(0.0157)
Boy	-0.331***	0.0275***	-0.0546***
	(0.0112)	(0.00369)	(0.00403)
Boy × Fiber	0.0165	0.00393	-0.00426
	(0.0238)	(0.00786)	(0.00804)
Observations	258,750	258,750	215,286
R-squared	0.669	0.571	0.574
Sample mean	.	0.10	0.89

This table presents OLS estimates for the effect of fiber on GPA, dropout and eligibility rates. Fiber coverage and GPA is measure at the year of (potential) graduation. Controls include sibling fixed effects, birth order indicators, a parish linear trend and municipality by year fixed effects. Standard errors clustered at the parish level. Our sample excludes movers and students in parishes with a negative reported change in fiber coverage between 2007 and 2011. *** p<0.01, ** p<0.05, * p<0.1

where D is an indicator for dropout or eligibility status for a student in family h, residing in parish p and (potentially) graduating in year t. Our sample consists of students graduating between 2007 and 2012 with at least one sibling graduating during the same period. As a fiber measure, we use fiber coverage in the year of graduation to increase sample size. We exclude families who move between the graduation of the graduation of the eldest and youngest sibling. In addition to a sibling fixed effect, we include indicators for birth order. Birth order has been shown to be an important determinant of educational attainment (Black et al., 2005). Since younger siblings will always be more intensely treated as coverage increases over time, birth order represents an important confounder in a sibling fixed effects specification.

As shown in table 7, we do not find any significant effects on either margin. Running the family fixed-effects regression for GPA reveals a point estimate similar to our baseline regression (3 percent of a standard deviation), but we cannot statistically distinguish the effect from zero. Since we do not include individual fixed effects, we can interact gender with treatment. Although boys have lower GPAs, lower eligibility rates and a higher risk of dropping out, we do not find evidence of a differential in the broadband effect in this specification. The sibling fixed effects specification suffers from a greater risk of bias due to unobservable individual heterogeneity, birth order effects and intra-family spillovers. In addition, as our last cohort graduates in 2012, we are missing a lot of younger siblings.

5 Conclusion

In this study, we provide evidence of a negative effect of high-speed broadband on the GPA of students in upper secondary school. Increasing local coverage from zero to 100 percent reduces expected GPA by almost 5 percent of a standard deviation in our preferred specification. The effect is larger for boys and for children born to parents with low education. A back-of-the-envelope LATE calculation suggests an effect almost twice that of our ITT estimate. We find that parental resources can mitigate the negative effect even when accounting for differences in student ability.

Our evidence suggests that part of an increased GPA gender gap can be explain by the fiber effect being more harmful to boys. Fortin et al. (2015) identify lowered post-secondary academic ambitions among boys as one of the main causes of a widening GPA gender gap. Our findings suggest that the GPA gap has increased due to boys reducing academic effort in favor of leisure, which is consistent with decreasing post-secondary ambitions and the labor market effects of leisure focused technology documented by Aguiar et al. (2017). At the extensive margin, we do not find any evidence suggesting an increased probability of dropping out of upper secondary school, although our identifying assumptions here are much stricter.

Our estimates should be interpreted as the effect of going from zero to full local coverage. While this may seem exaggerated, full coverage is not far from the explicit near-term goals stated by policy makers in Sweden and the EU (Regeringen, 2016). As countries approach full coverage, coming generations of students are less and less likely to have experienced a time without high-speed internet. While the effects we find could be driven by novelty, it is also true that online activities are rapidly evolving, making intertemporal comparisons difficult to interpret.

Estimating an effect of infrastructure abstracts from the fact that the effect is driven by the applications and services used by the student. This is rapidly evolving and notoriously hard to measure. Data on usage is often a closely held secret of service providers. While we provide some evidence that total time spent online has increased, we are unable to link this to local coverage. Thus, while we believe that student time use is the most likely mechanism, we cannot show that faster connections speeds causes students to spend more time online. Linking actual usage of online services to measures of student achievement seems like a worthwhile future research endeavor.

References

Aguiar, M., Bils, M., Charles, K. K., and Hurst, E. (2017). Leisure Luxuries and the Labor Supply of Young Men. Working Paper 23552, National Bureau of Economic Research. DOI: 10.3386/w23552.

Akerman, A., Gaarder, I., and Mogstad, M. (2015). The Skill Complementarity of Broadband Internet. *The Quarterly Journal of Economics*, 130(4):1781–1824.

Björklund, A. and Salvanes, K. G. (2011). Chapter 3 - Education and Family Background: Mechanisms and Policies. In Hanushek, E. A., Machin, S., and Woessmann, L., editors, *Handbook of the Economics of Education*, volume 3, pages 201–247. Elsevier. DOI: 10.1016/B978-0-444-53429-3.00003-X.

Black, S. E., Devereux, P. J., and Salvanes, K. G. (2005). The More the Merrier? The Effect of Family Size and Birth Order on Children's Education. *The Quarterly Journal of Economics*, 120(2):669–700.

Faber, B., Sanchis-Guarner, R., and Weinhardt, F. (2015). ICT and Education: Evidence from Student Home Addresses. Technical Report w21306, National Bureau of Economic Research, Cambridge, MA. DOI: 10.3386/w21306.

Fortin, N. M., Oreopoulos, P., and Phipps, S. (2015). Leaving Boys Behind: Gender Disparities in High Academic Achievement. *Journal of Human Resources*, 50(3):549–579.

Fudenberg, D. and Levine, D. K. (2006). A Dual-Self Model of Impulse Control. *American Economic Review*, 96(5):1449–1476.

Grover, S., Ensafi, R., and Feamster, N. (2016). A Case Study of Traffic Demand Response to Broadband Service-Plan Upgrades. In *Passive and Active Measurement*, Lecture Notes in Computer Science, pages 124–135.

Guldi, M. and Herbst, C. M. (2016). Offline effects of online connecting: the impact of broadband diffusion on teen fertility decisions. *Journal of Population Economics*, pages 1–23.

Hattie, J. (2015). Hattie effect size list - 256 Influences Related To Achievement.

Kenny, R. and Kenny, C. (2011). Superfast broadband: is it really worth a subsidy? *info*, 13(4):3–29.

Kolko, J. (2012). Broadband and local growth. *Journal of Urban Economics*, 71(1):100 – 113.

Malamud, O. and Pop-Eleches, C. (2011). Home Computer Use and the Development of Human Capital. *The Quarterly Journal of Economics*, 126(2):987–1027.

OECD (2011). National broadband plans. *OECD Digital Economy Papers*, 181.

OECD (2014). Does Homework Perpetuate Inequities in Education? `Retrivedfromhttps://www.oecd-ilibrary.org/content/paper/5jxrhqhtx2xt-en`.

Regeringen (2016). Sverige helt uppkopplat 2025. Retrived from https://www.regeringen.se/informationsmaterial/2016/12/sverige-helt-uppkopplat-2025---en-bredbandsstrategi/.

Schoenebeck, S. Y. (2014). Giving up Twitter for Lent: how and why we take breaks from social media. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, pages 773–782, Toronto, Ontario, Canada. ACM Press.

Statistics Sweden (2008). Use of computers and the internet by private persons in 2007. Retrived from https://www.scb.se/statistik/_publikationer/IT0102_2007A01_BR_IT01BR0701.pdf.

Statistics Sweden (2013). Use of computers and the Internet by private persons in 2012. Retrieved from https://www.scb.se/statistik/_publikationer/LE0108_2012A01_BR_IT01BR1301.pdf.

Svensson, M., Larsson, S., and De Kaminski, M. (2013). Professionalization, Gender and Anonymity in the Global File Sharing Community. *Piracy Effect*, pages 1–8.

Swedish Media Council (2013). Ungar och medier 2013. Retrived from https://statensmedierad.se/publikationer/ungarochmedier/ungarmedier201213.537.html.

Swedish Media Council (2015). Ungar och medier 2015. Retrived from http://statensmedierad.se/publikationer/ungarochmedier/ungarmedier2015.381.html.

Swedish Post and Telecom Authority (2008). Bredbandskartläggning 2007. Retrived from https://www.pts.se/upload/Rapporter/Internet/2008/Bredbandskartlaggning_2007_2008_5.pdf.

Thaler, R. H. and Shefrin, H. M. (1981). An Economic Theory of Self-Control. *Journal of Political Economy*, 89(2):392–406.

The Broadband Committee, SOU 2000:111 (2000). It-infrastruktur för stad och land. Retrived from http://www.regeringen.se/rattsdokument/statens-offentliga-utredningar/2000/11/sou-2000111-/.

The Swedish Internet Foundation (2013). Bredbandskollen – surfhastighet i sverige 2008- 2013.

Twenge, J. M., Martin, G. N., and Campbell, W. K. (2018). Decreases in psychological well-being among American adolescents after 2012 and links to screen time during the rise of smartphone technology. *Emotion (Washington, D.C.)*, 18(6):765–780.

Vigdor, J. L., Ladd, H. F., and Martinez, E. (2014). Scaling the digital divide: Home computer technology and student achievement. *Economic Inquiry*, 52(3):1103–1119.

A Appendix

Table A.1: Summary statistics by location information

	Full sample	With location	Without location
Year of birth	1991.84	1991.83	1991.86***
Boy	0.50	0.49†	0.68***
GPA at graduation	14.08	14.08	13.99***
GPA, year 1	13.75	13.76	13.68***
GPA, years 2-3	14.13	14.14	14.04***
GPA, 9th grade	222.93	223.01	221.15***
Fiber coverage, year 1	0.30	0.30	.
Fiber coverage, years 2-3	0.42	0.42	.
Academic program	0.47	0.47	0.49***
Immigrant	0.06	0.05†	0.12***
Mother's years of schooling	11.87	11.87	11.95***
Father's years of schooling	11.80	11.79	11.88***
Mother's income (log SEK)	11.46	11.46	11.41***
Father's income (log SEK)	11.86	11.86	11.82***
Number of students	261,101	250,433	10,668

*** Different from column 1 at the 1 percent level. † Different from column 1 at the 1 percent level.

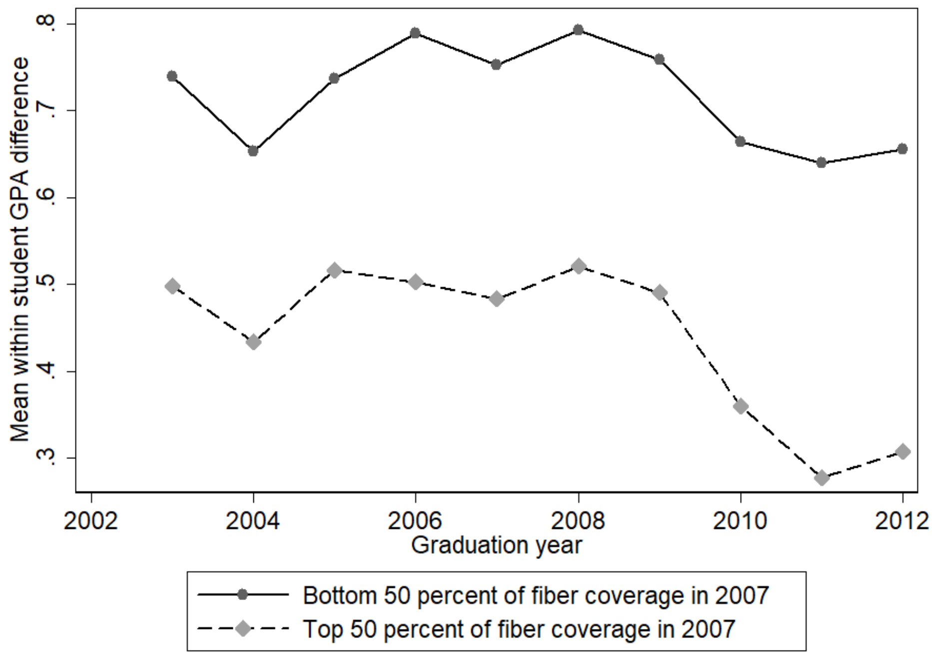

Figure A.1: Mean within student GPA difference by initial (2007) coverage.

Table A.2: Parish matching sequence

Match	Cumulative share of students with location data
Students at time t	38.5%
Parents at time t*	53.9%
Younger siblings at time t†	56.1%
Mothers at time t	79.1%
Fathers at time t	89.0%
Students at time t-1	89.5%
Mothers at time t-1	93.0%
Fathers at time t-1	94.5%
Mothers at time t+1	95.3%
Fathers at time t+1	95.8%

* Conditional on observing the same parish for both mother and father

† As students graduating in 2012 is the last cohort on record in our data, we are missing a lot of younger siblings.

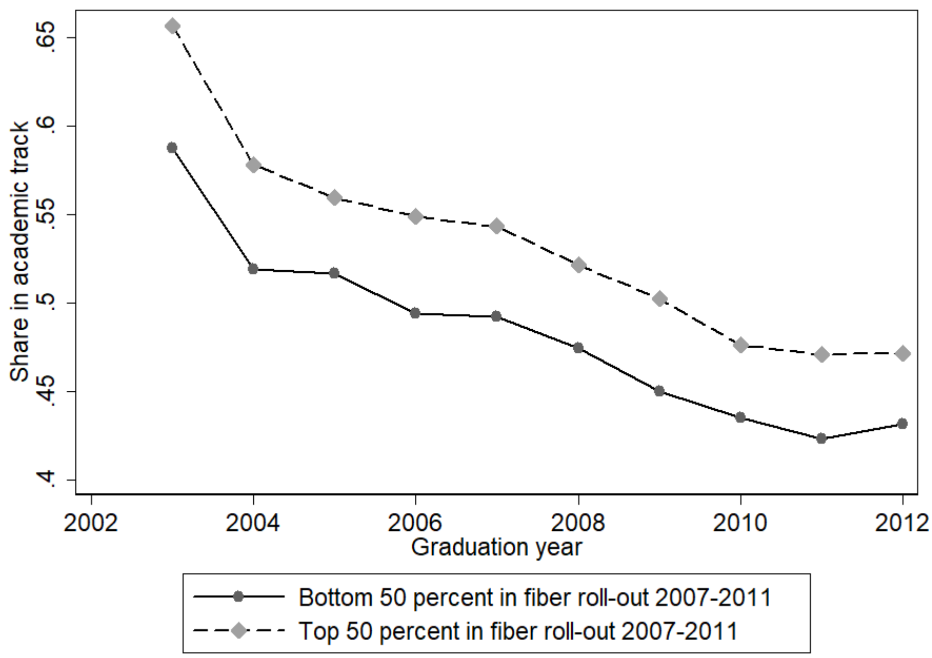

Figure A.2: Trends in the share of students in an academic program by treatment intensity