DISCUSSION PAPER SERIES

No. 11046

SPEED 2.0 - EVALUATING ACCESS TO UNIVERSAL DIGITAL HIGHWAYS

Gabriel Ahlfeldt, Pantelis Koutroumpis
and Tommaso Valletti

INDUSTRIAL ORGANIZATION and INTERNATIONAL TRADE AND REGIONAL ECONOMICS

ISSN 0265-8003

SPEED 2.0 - EVALUATING ACCESS TO UNIVERSAL DIGITAL HIGHWAYS

Gabriel Ahlfeldt, Pantelis Koutroumpis and Tommaso Valletti

Discussion Paper No. 11046
January 2016
Submitted 06 January 2016

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: (44 20) 7183 8801
www.cepr.org

SPEED 2.0 - EVALUATING ACCESS TO UNIVERSAL DIGITAL HIGHWAYS[†]

Abstract

This paper shows that having access to a fast Internet connection is an important determinant of capitalization effects in property markets. Our empirical strategy combines a boundary discontinuity design with controls for time-invariant effects and arbitrary macro-economic shocks at a very local level to identify the causal effect of broadband speed on property prices from variation that is plausibly exogenous. Applying this strategy to a micro data set from England between 1995 and 2010 we find a significantly positive effect, but diminishing returns to speed. Our results imply that disconnecting an average property from a high-speed first-generation broadband connection (offering Internet speed up to 8 Mbit/s) would depreciate its value by 2.8%. In contrast, upgrading such a property to a faster connection (offering speeds up to 24 Mbit/s) would increase its value by no more than 1%. We decompose this effect by income and urbanization, finding considerable heterogeneity. These estimates are used to evaluate proposed plans to deliver fast broadband universally. We find that increasing speed and connecting unserved households passes a cost-benefit test in urban and some suburban areas, while the case for universal delivery in rural areas is not as strong.

JEL Classification: H4, L1 and R2
Keywords: capitalization, digital speed, internet, property prices and universal access to broadband

Gabriel Ahlfeldt g.ahlfeldt@lse.ac.uk
London School of Economics and CEPR

Pantelis Koutroumpis p.koutroumpis@imperial.ac.uk
Imperial College London

Tommaso Valletti t.valletti@imperial.ac.uk
Imperial College London, University of "Tor Vergata" and CEPR

[†] We thank the Editor, five anonymous referees, as well as Kris Behrens, Brahim Boualam, Donald Davis, Gilles Duranton, Oliver Falck, Steve Gibbons, Shane Greenstein, Stephan Heblich, Christian Hilber, Hans Koster, Marco Manacorda, Jos van Ommeren, Henry Overman, Ignacio Palacios Huerta, Olmo Silva, Daniel Sturm, Maximilian von Ehrlich, and seminar participants in Barcelona, Bilbao, Boston (NBER Summer Institute), Florence, Kiel, London (SERC and Ofcom), Paris, Rome, St. Petersburg, Torino, Washington D.C. and Weimar for very useful comments.

1 Introduction

The importance of speed is well recognized. Higher speed brings workers and firms closer together and increases welfare due to travel-time savings and agglomeration benefits. Infrastructure projects—such as new metro lines, highways, high-speed rail or airports, all of which presumably increase speed within or between cities and regions—have long been popular among policy makers. The economic impact of such projects is well understood, and supportive evidence is relatively robust (see e.g. Baum-Snow, 2007; Baum-Snow and Kahn, 2000; Duranton et al., 2014; 2011; Faber, 2014).

In this paper, we deal with a different type of speed: digital speed. Does it matter how quickly one can surf the Internet using broadband? The possibilities that come with a faster Internet are countless: video streaming, e-commerce, or telecommuting, to name just a few. In a recent bestseller, Michael Lewis (2014) argues that superfast connections have even been used by high-frequency traders to rig the US equity market.[4] In contrast to the classic infrastructures mentioned above, it is normally left to the market to supply Internet connections, via Internet Service Providers such as telecom and cable providers. Policy makers have traditionally limited their interventions to a few targeted rural areas. Perhaps as a way to escape the economic crisis, this discreet approach has changed recently. In the US, the Federal Communications Commission (FCC) launched the National Broadband Plan in 2010 to improve Internet access. One goal is to provide 100 million American households with access to 100 Mbit/s connections by 2020.[5] In Europe, broadband is one of the pillars of Europe 2020, a ten-year strategy proposed by the European Commission. Its Digital Agenda identifies targets that are as aspiring as the US's: also by 2020, *every* European citizen will need access to at least 30 Mbit/s.[6]

We argue that it is possible to infer the value brought by a faster Internet connection via changes in property prices. Theoretically, it is evident that fixed broadband, by far the usual way people connect to the fast Internet, comes bundled with a property whose price might, therefore, be affected. Broadband availability and speed embody just one characteristic of a property that contributes to determining its value (along with local amenities, infrastructure, and other neighborhood characteristics). Anecdotal evidence makes a strong case that broadband access is an important determinant of capitalization effects in property markets. In 2012, *The Daily Telegraph*, a major UK daily newspaper, reported the results of a survey among 2,000

[4] Using fiber-optic cables that link superfast computers to brokers, the high-frequency traders intercepted and bought the orders of some stock traders, selling the shares back to them at a higher price and pocketing the margin. The key to this scheme was an 827-mile cable running from Chicago to New Jersey that reduced the journey of data from 17 to 13 milliseconds (Lewis, 2014).

[5] http://www.broadband.gov/plan.

[6] Additionally, at least 50% of European households should have Internet connections above 100 Mbit/s; see http://ec.europa.eu/digital-agenda/our-goals/pillar-iv-fast-and-ultra-fast-internet-access.

homeowners, showing that a fast connection is one of the most important factors sought by prospective buyers. The article states that "[...] a good connection speed can add 5 percent to a property's value." Perhaps more tellingly, the survey says that one in ten potential buyers reject a potential new home because of a poor connection, and that, while 54% considered broadband speed before moving in, only 37% looked at the local crime rate.[7] Rightmove, one of the main online real estate portals in the UK, rolled out a new service in 2013 to enable house hunters to discover the broadband speed available at any property listed on the site, along with more-typical neighborhood information such as transport facilities or schools.[8]

To empirically estimate the valuation for broadband speed via the variation in house prices, we have access to very detailed information about broadband development and residential properties for the whole of England, over a rather long period (1995-2010). We find an elasticity of property prices with respect to speed of about 3% at the mean of the Internet speed distribution. However we also find diminishing returns—that is, the increase in value is greater when starting from relatively slow connections, which helps to put the empirical results in the right perspective. The average property price increased by 2.8% when going from a slow narrowband dial-up connection to the first generation of ADSL broadband Internet connections, which allowed a speed of up to 8 Mbit/s. The price increased by an additional 1% when a newer technology, ADSL2+, was rolled out to offer Internet speeds up to 24 Mbit/s. In other words, families are willing to pay a premium of 1% of the property price, or about £2,200 (≈$3,300) when, other things equal, the property is supplied by a fast connection compared to a normal broadband connection. This effect corresponds to an increase in school quality by one third of a standard deviation (Gibbons et al., 2013) or a reduction in distance to the nearest London underground station of one third of a kilometer (Gibbons and Machin, 2005). The magnitude of the effect is smaller than, e.g., the negative effect of having a convicted sex-offender living nearby (4%, see Linden and Rockoff, 2008) or the positive effect of a good grade awarded to the local school in a school quality review (8.7%, Figlio and Lucas, 2004), but more sizable than the effect of the clean-up of a hazardous waste site (Greenstone and Gallagher, 2008).

We further decompose these average results by income and degree of urbanization. It turns out that the gains are very heterogeneous, and they are highest at the top of the distribution, among the richest people living in the most densely populated areas, London in particular. Put differently, these results imply that, on average, a household would be willing to spend, over and above the subscription fee to the Internet provider, an extra £8 (≈$12) per month for the option

---

[7] http://www.telegraph.co.uk/property/propertynews/9570756/Fast-broadband-more-important-to-house-buyers-than-parking.html.

[8] http://www.rightmove.co.uk/broadband-speed-in-my-area.html. Prior to this service, people looked for postcode-level speed information in broadband provider websites, forum discussions, and web-based speed checkers. This type of information started to appear with the launch of the first ADSL connections in the early 2000s; see, e.g., http://forums.digitalspy.co.uk/showthread.php?t=190825.

to connect to the high speed ensured by ADSL2+ compared to an otherwise identical property that only had access to a basic ADSL connection. In rich and dense places like London the surplus can be as high as £25 (≈$37.5) per month. Endowed with these findings, we then evaluate the benefits of the EU Digital Targets for different regions in England, which we compare with available cost estimates. We find that increasing speed and connecting unserved households passes a cost-benefit test in urban areas, while the case for universal delivery in rural areas is not very strong.

In order to provide reliable estimates of the valuation for broadband speed, we need to avoid the circular problem that is present in all spatial concentrations of economic activities. First, we need to separate the effect of high broadband speed on property prices from other favorable locational characteristics, such as good transport access or schools. Second, the available speed is endogenous to factors that determine broadband demand and are likely correlated with property prices, such as high levels of income and education levels. Thus, to avoid spurious correlation, we have to account for macroeconomic shocks such as gentrification that potentially affect speed and property prices simultaneously.

We are able to trace the presence of broadband, and its speed, at the level of each local delivery point, called a Local Exchange (LE) in the UK (this would be called the Central Office in the US). Every home can be supplied by one and only one LE, which we can perfectly identify. Within a given LE area, the distance between the user's premises and the LE is, by far, the most important factor affecting the performance of a given connection. In addition, LEs have been upgraded at different points in time, with some exchanges boasting faster technologies than others. The local distribution from legacy phone networks does not influence phone quality but does affect broadband quality. This provides us with an ideal variation of speed over time within an extremely small area. We are able to identify the causal effect of digital speed on property prices from two alternative sources of variation. First, we exploit a discontinuity *across* LE boundaries over time. Adjacent properties can belong to the catchment areas of different LEs and, therefore, with different distances to the exchange and possibly also different vintages of technology. Holding constant all shocks to a spatially narrow area along the boundary of two LEs, the discontinuous changes in speed that arise from LE upgrades at both sides of such a boundary provide variation that is as good as random. In other words, we compare the house prices of two properties, located next to each other, that are observationally equivalent in terms of characteristics but for the speed available to each one of them. Second, we use variation over time *within* LEs. Because we can hold constant any macroeconomic shock that mutually determines property prices and upgrade decisions, which are made at the LE level, the conditional variation in speed is plausibly exogenous. Both identification strategies result in very similar estimates.

Our work is related to two streams in the literature. In general, our methods are common to a large literature in urban and public economics that has explored capitalization effects of local public goods or non-marketed externalities more generally (Chay and Greenstone, 2005; Davis, 2004; Greenstone and Gallagher, 2008; Linden and Rockoff, 2008). We use similar methods and show how they also can be used in settings where, a priori, one would not think of an externality. Here, we deal with a market that is largely competitive and privately supplied, but there are still capitalization effects: a good part of the consumer surplus associated with broadband consumption seems to go to the property seller as a scarcity rent, and not to the broadband suppliers.

A second stream in the literature to which we contribute is related to the evaluation of broadband demand and of the benefits associated with Internet deployment. At a macro level, Czernich et al. (2011), using a panel of OECD countries, estimate a positive effect that Internet infrastructure has on economic growth. Kolko (2012) also finds a positive relationship between broadband expansion and local growth with US data, while Forman et al. (2012) study whether the Internet affects regional wage inequality. Greenstein and McDevitt (2011) provide benchmark estimates of the economic value created by broadband Internet in the US. Some studies assess the demand for residential broadband: Goolsbee and Klenow (2006) use survey data on individuals' earnings and time spent on the Internet, while Nevo et al. (2015) employ high-frequency broadband usage data from one ISP. To our knowledge, ours is the first study to estimate consumer surplus from Internet usage using property prices for a large economy.

The rest of the paper is organized as follows. In Section 2, we describe the development of broadband Internet in England and discuss the theoretical linkage between broadband speed and property prices. Section 3 presents the empirical strategy. The main results are shown and discussed in Section 4. Section 5 uses the empirical findings to quantify the benefits for the EU 2020 digital targets. Finally, Section 6 concludes.

2 The broadband market

In this section, we first describe the recent development of broadband Internet in England and then give an overview of its variation over time and space. We then describe our data sources. Finally, we provide a simple theoretical model that links broadband availability, and its speed, to property prices.

2.1 The broadband market in England

The market for Internet services in England[9] is characterized by the presence of a network, originally deployed by British Telecom (BT) during the first part of the 20th century to provide voice telephony services. BT was state-owned until its privatization in 1984. This network consists of 3,897 Local Exchanges (LEs). Each LE is a node of BT's local distribution network (sometimes called the "local loop") and is the physical building used to house internal plant and equipment. From the LE, lines are then further distributed locally, by means of copper cables, to each building in which customers live or work, which tend to be within two kilometers from the LE. LEs aggregate local traffic and then connect up to the network's higher levels (e.g., the backbone) to ensure world-wide connectivity, typically by means of high-capacity (fiber) lines.

While the basic topology of BT's network was decided several decades ago, technology has proven extremely flexible. The old copper technology, until the end of the 90s, provided a speed up to 64 kbit/s per channel via dial-up (modem) connections. Without having to change the cables in the local loop, it has been possible to supply high-speed Internet by installing special equipment in the LEs. A breakthrough occurred with a family of technologies called DSL (Digital Subscriber Line), which use a wider range of frequencies over the copper line, thus reaching higher speeds. The first major upgrade program involved bringing the ADSL technology to each LE. BT began the program in early 2000 and took several years to complete it. This upgrade could initially improve Internet speed by a factor 40 compared to a standard dial-up modem and, afterwards, allowed speeds up to 8 Mbit/s.

Along with technological progress, the regulatory framework and the competitive landscape also evolved over the same period. Ofcom, the UK's regulator for communications, required BT to allow potential entrants to access its network via the so-called "local loop unbundling" (LLU). LLU is the process whereby BT makes its local network of LEs available to other companies. Entrants are then able to place their own equipment in the LE and to offer services directly to customers. LLU started to gain pace in 2005, and entrants have progressively targeted those LEs located in more densely populated areas. Regulatory intervention is limited to wholesale prices, while retail prices are freely set by competing providers. A further major improvement occurred with ADSL2+. This upgrade, which allows for download speeds, theoretically, up to 24Mbit/s, started around 2007. It was first adopted by some of the new LLU entrants, and BT followed with some lag. ADSL, LLU, and ADSL2+ are going to be major shifters of speed in our data, as they varied substantially over time and by LE. In addition, all technologies based on DSL are "distance-sensitive" because their performance decreases significantly as you get further away from the relevant LE.

[9] The broadband description applies to the whole of the UK. However, since our property data cover only England, we always refer to England alone throughout the paper.

Of course, the diffusion of broadband internet was not uniform across the UK, and several demand and supply factors determined different penetration rates across markets and over time. Nardotto et al. (2015) document how the entry process took off around 2005, and show that entrants improved considerably the speed available locally in each LE where they entered. First, local entry of new providers was the main reason for the adoption of broadband internet. In order to recover entry's large investment, entrants first unbundled the larger and more profitable LE-markets, and later expanded to cover a large share of the country. Second, the shape and the size of the area covered by each LE was an important determinant of entrants' costs. Finally, rapid technological progress, along with entrants' learning curves, decreased costs over time.[10]

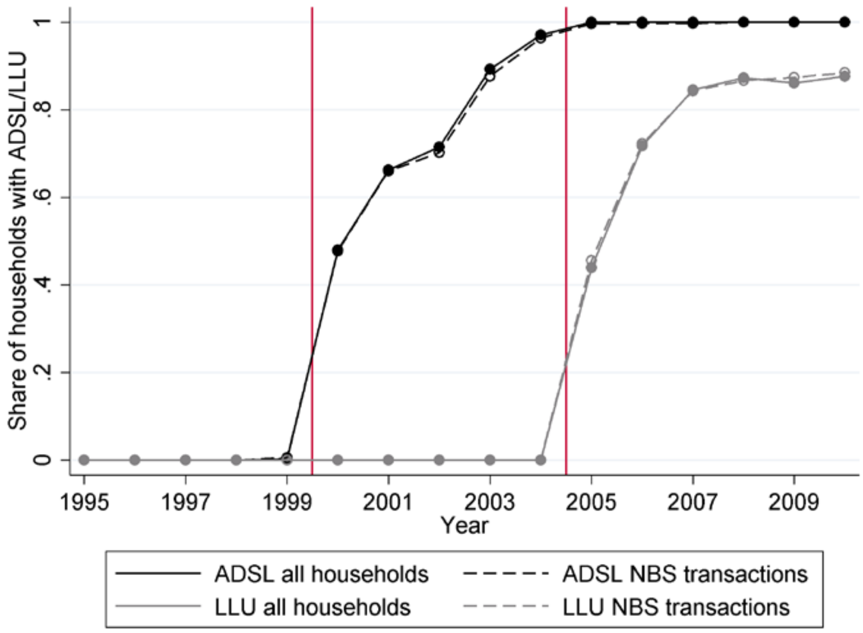

Notes: Black (grey) lines refer to ADSL (LLU) activation. Solid (dashed) lines refer to all households in England (NBS = Nationwide Building Society transactions data set)

Figure 1: Share of households with ADSL/LLU over time

Figure 1 shows the evolution of the share of English households in the catchment area of LEs enabled with ADSL (black solid line) or with LLU entrants (grey solid line).[11] We therefore cover the period that was crucial for the development of residential Internet. The share of properties in our sample reflects very closely the technological pattern in England (dashed lines), providing reassurance on its representativeness. In Appendix A, we provide further empirical evidence, showing maps of how these technological changes occurred by region and over time.

[10] See Chen and Savage (2010) for a related analysis for the US.
[11] We do not show ADSL2+ in order not to clutter the figure, but it would lie below the LLU curve.

Figure 2 is a static map of a few Local Exchanges located north of London. The figure reports the location of the relevant LEs in that area (big black dots), and their catchment areas, based on the full postcodes served (black boundaries). Each colored dot represents the location of one transaction in the property dataset, where lighter colors correspond to increasing distances from the exchange (from red to yellow). Black icons denote groups of properties that have been matched to common boundary segments. These two figures show two important things that will inform our empirical strategy. First, there is considerable variation both in the distance between premises and the relevant LE (figure 2), and in the technology available over time at a given LE, which should have an impact on the available speed for a specific property (figure 1). We will, thus, be able to control for unobserved shocks to neighborhoods at very disaggregated levels and restrict identification to variation that stems from changes in the relative distribution of speeds within LEs over time. Second, there are enough properties at the LE boundary allowing us to exploit discontinuities in speed increases if one or both LEs are upgraded.

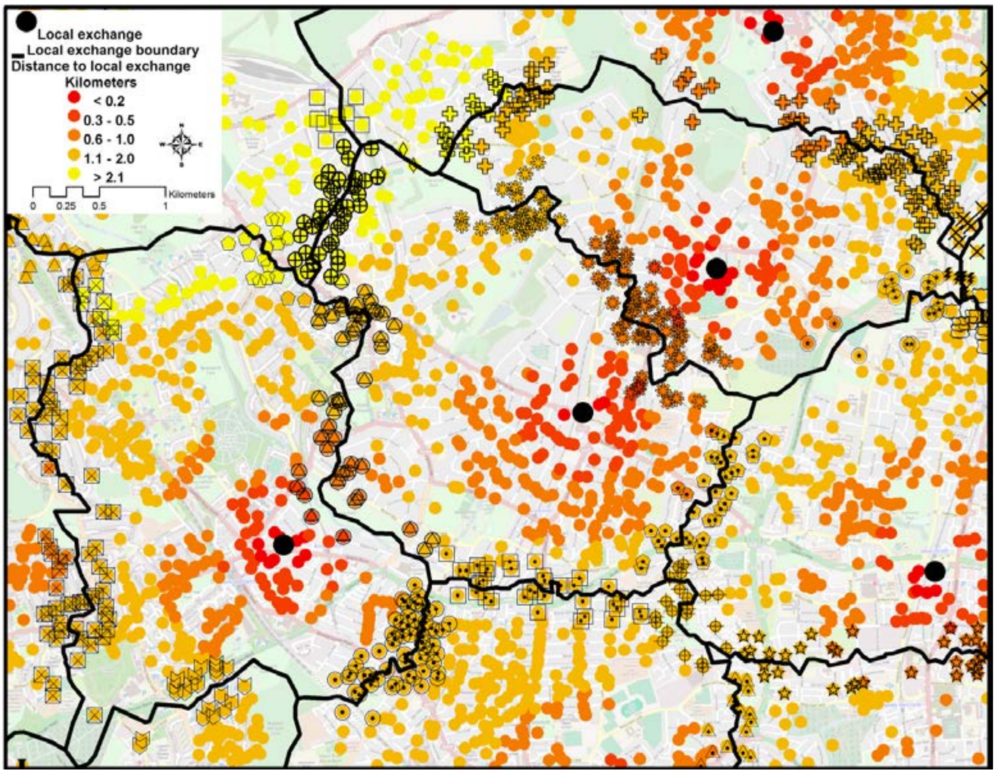

Notes: Black icons denote groups of properties within 200m of a shared boundary segment. The colored dots are transactions from the NBS dataset. The black dots are the locations of LEs and the black boundaries are their catchment areas, both from the Ofcom dataset.

Figure 2: Distribution of properties and LE catchment areas

To complete the picture, broadband Internet can also be supplied via an alternative cable network.[12] The cable operator Virgin Media deployed its own network during the 1990s,

[12] At the beginning of 2010, BT had a retail market share of 28 percent, the cable operator had a market share of approximately 22 percent, and the entrants (the main ones are TalkTalk, Sky, O2 and Orange) had the remaining 50 percent of the market. There has been little investment in fiber within the local loop, and during

primarily for the purpose of selling cable TV. The topology of this network is very different from BT's. It covers roughly 50% of the premises in England, concentrating its presence in urban areas and flat parts of the country. The cable network can be upgraded to support broadband only if an area is already covered by cable, which has not expanded its reach since the 1990s. Cable technology, since it also aims at providing TV, is typically faster than ADSL, and broadband speed does not degrade substantially with distance from the exchange.

2.2 Raw data

Our dataset stems from several sources. The main block concerns the development of broadband in England over the period 1995-2010. Ofcom has made available to us all the information it collects on the broadband market for regulatory purposes. The dataset comprises quarterly information at the level of each of the 3,897 LEs in England. For each local exchange, we know the precise coverage of BT's local network—that is, all the specific full postcodes served by a certain LE—and, therefore, we know how many buildings and total lines can eventually have broadband. We remark that a full postcode unit contains about 10-15 households, which are all connected to the same LE.[13]

We can identify when a LE was upgraded to ADSL or ADSL2+, and if and when it attracted entrants via LLU. We also know, in the catchment area of the LE, whether or not cable is available. Finally, we know how broadband penetration varies over time in a given LE, as we are told the total number of subscribers (via BT, via an entrant, or via cable), which can be compared to the total lines available locally to compute broadband penetration.

This detailed information was supplemented with information on broadband speed tests carried out by individuals in 2009 and 2010. We obtained three million tests from a private company.[14] For each individual/speed test, we observe the operator, the contract option chosen by the user, the location (full post code), as well as when the test was carried out. Thus, we can calculate the distance between the user's premises and the exact location of the relevant LE. The dataset contemplates two measures of performance: download speed and upload speed. We focus on the former, which is, by far, the more important feature for residential household users. It is important to note that, throughout the whole paper, we refer to the speed measured in the

the period we consider here, there has been limited take-up of high-speed connections based on 3G cellular technology. Broadband access via Wi-Fi technologies, on the other hand, is included in our dataset.

[13] A full (typically, 7 digit) postcode in the UK captures a narrowly defined area. There are approximately 2 million postcodes in the UK. A full postcode is not an address, but still covers areas that are on average within a radius of 50m, which gets even narrower in densely populated areas (e.g., 20m in London).

[14] http://www.broadbandspeedchecker.co.uk. More information is provided in Section 3.1.

dataset on speed tests as “actual” speed. This is not the same as the speed typically advertised by operators in their plans, to which we refer as “nominal” speed.[15]

For the analysis of the capitalization effects of broadband capacity, we use transactions data related to mortgages granted by the Nationwide Building Society between 1995 and 2010. The data for England comprise more than one million observations,[16] and include the price paid for individual housing units along with detailed property characteristics. These characteristics include floor space (m^2), the type of property (detached, semi-detached, flat, bungalow or terraced), the date of construction, the number of bedrooms and bathrooms, garage or parking facilities, and the type of heating. There is also some buyer information, including the type of mortgage (freehold or leasehold) and whether they are first-time buyers. Note that the transaction data include the full UK postcode of the property sold, allowing it to be assigned to grid-reference coordinates.

With this information, it is possible with GIS software to calculate distances to LEs. Furthermore, it is possible to calculate distances and other spatial measures (e.g., densities) for the amenities and environmental characteristics such as National Parks, as well as natural features such as lakes, rivers and coastline. The postcode reference also allows a merger of transactions and various household characteristics (median income and ethnic composition) from the UK census; natural land cover and land use; and various amenities, such as access to employment opportunities, retail services, cultural and entertainment establishments, school quality, and measures of online services (e.g., Amazon evening delivery, Uber fleet services). A more-detailed description of all the data used is in Appendix B. In Appendix C1, we also show that the distributions of other observable amenities do not differ discontinuously on the two sides of a LE.

2.3 A simple conceptual model

Unlike local public goods such as good (public) schools, public safety, or air quality, which are often analyzed in the house price capitalization literature, households subscribed to broadband pay a price to their Internet provider. A capitalization effect of broadband is, therefore, not an obvious feature of the spatial equilibrium. The purpose of this section is to introduce a simple model that links broadband speed to property prices. Our intention is not to introduce a model for structural estimation, but, rather, to think about this link in a simple and transparent manner.

---

[15] The discrepancy for the top plans is large and amounts to a factor 4 (results are available on request from the authors). This factor is also in line with independent findings of Ofcom; see, e.g., http://stakeholders.ofcom.org.uk/market-data-research/other/telecoms-research/broadband-speeds/speeds-nov-dec-2010/, and Figure 1.2 in particular).

[16] This represents 10% of all mortgages issued in England over the period.

For this purpose, imagine that there are n areas, indexed by j = 1, …, n. In each area there is a population of household buyers whose total number is normalized to unity. The value of a property is denoted as V, which can be made dependent on all its characteristics, such as number of rooms, local amenities, etc., except for broadband availability, which is described next. The price of a property is denoted as P.

Households are heterogeneous in their value of using broadband. Value can derive from different sources—from leisure (surfing the Internet) to being able to work from home. We are not interested in the particular channel, but simply imagine that people are heterogeneous in the way that they use and value the Internet. Let $v \cdot \log(q_j)$ denote the gross utility of household type v using a broadband of quality q_j, where q_j is the Internet quality available in area j, for instance, the speed of the connection. This specification reflects diminishing marginal returns to speed, as well as the fact that everybody would enjoy faster connections, ceteris paribus, despite heterogeneity in tastes. The distribution of household types v is assumed to be uniform between 0 and a_j in area j, thus the density is $1/a_j$.[17]

The consumers' choice is whether or not to purchase broadband, conditional on having bought a property. We normalize the payoffs from not using broadband to zero. Broadband of quality q_j is sold at a price p_j. Since broadband is a durable good, all these variables are to be interpreted as flows in each period. We also assume that, at some period in the future denoted as T_j, some alternative technology that does not need fixed lines becomes available, and it will be preferred by all customers (because it is cheaper or better, or both). Think, for instance, of LTE mobile technology replacing fixed broadband. The key point is that this technology will *not* be bundled with the property anymore, but it will represent a completely separate purchase that has nothing to do with a property. The cumulative utility for type v from fixed broadband access is thus $[v\log(q_j) - p_j]\Delta_j$, where $\Delta_j = \int_0^{T_j} e^{-\rho t} \mathrm{d}t = \frac{1-e^{-\rho T_j}}{\rho}$ and ρ is the discount rate. Note that, if the alternative technology never becomes available, $T_j \to \infty$ and the discount factor Δ_j simplifies to $1/\rho$, i.e., the value of a perpetuity.

Households whose value of broadband is high enough will purchase a broadband connection. In particular, the marginal broadband household in area j is defined by $v_j^* = p_j/\log(q_j)$, and all types between v_j^* and a_j purchase broadband in that area in every period.

On the property supply side, we assume that homes in a given area are scarce, such that sellers can always extract all buyers' net surplus. Alternatively, one can also assume that sellers are able to observe buyers' types—during negotiations, for example—and make take-it-or-leave-it offers

[17] The example is generalizable to a more general distribution function $F(v)$ that satisfies the monotone hazard rate condition.

leading to the same outcome. Households are assumed to be perfectly mobile, with reservation utility U. In a spatial equilibrium, house prices in area j will, therefore, be

$$P_j = \begin{cases} V - U \text{ for } v < v_j{}^* \text{ (households without broadband),} \\ V - U + [v\log(q_j) - p_j]\Delta_j \text{ for } v \geq v_j{}^* \text{ (households with broadband).} \end{cases} \tag{1}$$

Since the econometrician will not observe types, but just the average prices in a given area with or without broadband subscription, we can calculate these averages from (1) as

$$\bar{P}_j = \frac{(V-U)v_j{}^*}{a_j} + \int_{v_j{}^*}^{a_i} \frac{\left[V - U + \left(v\log(q_j) - p_j\right)\Delta_j\right]}{a_j} \mathrm{d}v = V - U + K_j\Delta_j, \tag{2}$$

where

$$K_j \equiv \frac{(a_j q_j - p_j)[(a_j q_j + p_j)\log(q_j) - 2q_j p_j]}{2a_j q_j^2}.$$

It is a matter of simple maths to show that[18]

a) $\frac{\partial \bar{P}_j}{\partial q_j} = \Delta_j \frac{\partial K_j}{\partial q_j} > 0$ and $\frac{\partial \bar{P}_j{}^2}{\partial q_j^2} < 0$,

b) $\frac{\partial \bar{P}_j}{\partial a_j} = \Delta_j \frac{\partial K_j}{\partial a_j} > 0$,

c) $\frac{\partial \bar{P}_j}{\partial T_j} = K_j \frac{\partial \Delta_j}{\partial T_j} > 0$.

Eq. (2) and the associated comparative statics confirm the intuition that broadband speed gets capitalized into house prices. In particular (part a), prices should increase in those areas with higher available speed q_j, and they increase at a decreasing rate (decreasing returns to speed). Prices should also increase (part b) in those areas where there is a higher willingness-to-pay for the Internet, because of the heterogeneity in the population that we have modelled via a_j (which may be related to income, something we do observe at the level of an area in our data). Places with the highest price premium for speed are likely also to have residents with the greatest taste for speed. Eq. (2) also makes a point about sorting: the coefficient estimates from the hedonic price regressions that we will run should return the mean marginal valuations of properties (Bayer et al., 2007), and one needs to be careful when conducting policy evaluation involving levels of speed different from those observed. Finally (part c), the capitalization effect depends on whether there is an expectation that fixed line broadband will be displaced by technologies that are not bundled with the property. If these technologies do not exist, our results effectively

[18] One just needs that $p_j < a_j\log(q_j)$, which must hold true for the problem to make economic sense, otherwise, not even the household with the highest willingness to pay would get a broadband subscription.

capture a perpetuity in the value of broadband, else they will capture only the net present value from a shorter period.

The model also has an ancillary prediction about broadband penetration in a given area. This provides a useful check for the robustness of our main results and a way to evaluate the channels through which the capitalization effect operates. Penetration is given by

$$Penetration = 1 - \frac{{v_j}^*}{a_j} = 1 - \frac{p_j}{a_j \log(q_j)}, \qquad (3)$$

which is also increasing in speed q_j, and at a decreasing rate. Eq. (3) also says that – ceteris paribus – penetration in a certain area is driven by Internet characteristics (q_j and p_j) and by population characteristics (a_j), but not amenities that depend themselves on the availability of fast broadband (e.g., cybercafés).

Note that we left the broadband subscription price p_j unmodelled, thus the main prediction that property prices increase with speed is independent of the precise market structure of the local broadband market: intuitively, it is stronger when the broadband supply is very competitive, but it holds even for a monopolist provider. In other words, there are limits to the consumer surplus that ISPs can appropriate when speed increases. Competition is the upper limit, in fact broadband subscription fees cannot increase with willingness to pay for speed when competition is intense, as they will just reflect costs. But even a monopolist would be constrained by its inability to observe different types perfectly and would, therefore, leave some information rent to higher types. Our approach presumes that all remaining consumer surplus from broadband, over and above the broadband price paid to the provider, is appropriated by the seller of the property. If this were not the case, then the impact that broadband might have on property prices would underestimate the consumer surplus from broadband use. We will return to this point in our conclusions.

3 Empirical framework

The primary aim of our empirical strategy is to provide a causal estimate of the impact of high-speed broadband supply on house prices. The empirical challenge in estimating this causal effect is to separate the effect of broadband supply from other unobserved and potentially correlated determinants of house prices. In particular, we must ensure that there are no omitted variables that simultaneously determine broadband supply and house prices. We argue that robust identification can be achieved from discontinuous variation in speed *over* time and *across* LE boundaries. Variation over time helps disentangle the effect of broadband supply from unobserved (spatially) correlated locational factors, such as good transport access or better schools. By further placing properties into groups that are near to and share the same LE boundary, it is possible to control for shocks at a very small spatial level. We argue that variation in speed over time across a LE boundary within such a small area is plausibly exogenous and as good as random. We also run an alternative identification which relies on the comparison of house prices to broadband supply *over* time and *within* LE areas. Decisions that affect the broadband supply of a property are generally taken at the level of the LE serving an area. Conditional on shocks to a certain LE catchment area—such as a sudden increase in income or education of the local population—within-LE variation in speed over time that results from the distance of a property from the relevant exchange can be assumed to be exogenous.[19]

We follow the popular hedonic pricing method to separate various determinants of property prices. Rosen (1974) has provided the micro-foundations for interpreting parameters estimated in a multivariate regression of the price of the composite housing good against several internal and locational characteristics as hedonic implicit attribute prices. Underlying the hedonic framework is the idea that, given free mobility in spatial equilibrium, all locational (dis)advantages must be offset by means of property price capitalization. There is a long tradition in the literature—dating back at least as far as Oats (1969)—that made use of the hedonic method to value local public goods while holding confounding factors constant. One of the typical challenges faced by such hedonic valuation studies is the potential for bias due to omitted variables that are correlated with a phenomenon of interest. Recent applications of the hedonic method have tackled this problem by making use of variation over time to identify the effects of locational improvements from unobserved time-invariant locational factors (Ahlfeldt and Kavetsos, 2014; Chay and Greenstone, 2005; Davis, 2004; Linden and Rockoff, 2008).

Both of the empirical specifications we employ are drawn from this line of research. We model the (log) price of a property sold at a full postcode *i* at time *t*, served by LE *j* and lying on the LE boundary segment *k* as a function of the available broadband speed, as well as a range of internal and locational property characteristics that are partially observed and partially unobserved. Our baseline empirical specification is a variant of a spatial boundary discontinuity design (BDD):

$$\log(P_{ijkt}) = \sum_{m=1}^{2} \alpha_m (S_{ijt})^m + \sum_{n=1}^{4} \tau_n (DIST_{ij})^n + \mathrm{X}_i' \mu_t + \psi_{kt} + \varphi_j + \epsilon_{ijt}, \qquad (4)$$

where S_{ijt} is the available broadband speed, and $DIST_{ij}$ is the Euclidian distance from a postcode *i* to the relevant LE *j*. We use a quadratic specification for broadband speed to allow the property price to vary non-linearly with speed, as predicted by our simple model. The distance

[19] Note that local exchange areas are relatively small. The median radius of a local exchange area is less than six km, as far as old voice telephony services are concerned. As for broadband, the area where it can be supplied effectively is smaller, up to 2-3 km from the local exchange, as shown below in the results. In cities, the median radius of a LE is further reduced—e.g., less than 2 km in London.

polynomial controls for unobserved time-invariant locational characteristics that are correlated with distance to the LE. As discussed in more detail in the next section, our variable of interest S_{ijt} is constructed using fourth-order polynomials of $DIST_{ij}$ following an engineering literature. Because S_{ijt} varies over time, the speed effect, after controlling for the time-invariant distance trend, is identified from variation over time. The control variable approach is therefore equivalent to postcode fixed effects in terms of its power to absorb unobserved locational effects that are correlated with S_{ijt}. Compared to the alternative of using postcode fixed effects, we prefer this control variable approach because of a relatively limited number of repeated sales at the same postcode level.[20] X_i' is a vector of property and locational characteristics discussed in the data section. This is interacted with a full set of year effects, so that μ_t is a matrix of implicit prices for attribute-year combinations. φ_j is a dummy to control for unobserved time-invariant LE effects. Finally, k indexes properties that lie along the same boundary segment that separates two LE areas. We match properties in LE j to the nearest property in LE $l \neq j$ and define a common time-varying fixed effect ψ_{kt} for properties in j whose nearest neighbor is in l and vice versa. These fixed effects ensure that we identify from a differential increase in speed at the two sides of the boundaries, holding constant all other time-varying effects that are common to both sides of a boundary. Figure 2 illustrates the matching of properties across adjacent LEs.

This specification exploits the discontinuity at the boundaries between LEs. Overall, there are 86,569 LE boundary x year effects in our data, which denote boundary segments that are common to the same two LEs. With this specification, we attribute differences in price changes over time across a common boundary to the respective differences in speed changes over time. We restrict our sample to properties that are close to a LE boundary to explicitly exploit the spatial discontinuities in speed changes that arise across a LE boundary if the broadband infrastructure is altered. We note that a discontinuity arises not only if just one of two adjacent LEs is upgraded, but also if both LEs are upgraded, and the distance to the respective LEs differs significantly at both sides of the LE boundary. Because, at a local level, the allocation of a property to either side of the same boundary is as good as random, it is unlikely that unobserved shocks affect speed and property prices on one side of the boundary but not on the other. Even in this unlikely event, such shocks are absorbed by the LE boundary x year effects.

We also estimate an alternative specification in which we replace the LE boundary x year effects with a set of 37,804 LE x year fixed effects φ_{jt} that control for all macroeconomic shocks at the LE level:

[20] Less than half (15 percent) of the full postcodes in the Nationwide data set contain two (three) or more transactions. On average, there are 2.15 transactions per full postcode over the 15-year period we cover.

$$\log(P_{ijt}) = \sum\nolimits_{m=1}^{2} \alpha_m (S_{ijt})^m + \sum\nolimits_{n=1}^{4} \tau_n (DIST_{ij})^n + \mathrm{X}_i' \mu_t + \varphi_{jt} + \epsilon_{ijt}, \quad (5)$$

With this specification we focus on a different source of variation, compared to eq. (4). Instead of exploiting discontinuous variation in speed over time across LE boundaries we now identify exclusively from continuous variation in speed over time within LEs. In estimating eq. (5) we also use the universe of transactions and variation in speed, which helps addressing the external validity problem inherent to all boundary discontinuity designs. This specification delivers a causal effect of broadband speed on house prices under the identifying assumption that year-specific shocks that potentially determine broadband capacity are uncorrelated with distance to the LE within the area that the LE serves. This is a plausible assumption for two reasons. First, any change to the LE technology will affect the entire catchment area served by the LE, so it is rational for broadband suppliers to base decisions on the average trend in this area. It is, therefore, unlikely that within-LE shocks that might affect property prices—e.g., an income increase among the population near the LE relative to other areas—would also affect the technological upgrading decisions above and beyond their effect on the LE area average, which is captured by φ_{jt}. Second, LEs serve relatively small areas, with a layout that was defined decades ago and boundaries that do not line up with spatial statistical units, such as census wards. The catchment area of each LE is typically known only to providers and is not used to create any other related boundaries. Reliable information on year-on-year changes at the sub-LE area level is difficult to obtain, which makes it unlikely that providers would be able respond to within LE-area shocks even if they wanted to.[21] This specification is arguably more open to criticism because there may be within-LE trends in property prices that are correlated with distance to the LE, something that is absent with the previous specification relying on the boundary discontinuity. It is noteworthy that the interactions of year effects and attributes X_i' flexibly control for property price trends that are correlated with any of the observable structural and locational characteristics. Conditional on these controls, it is less likely that within-LE trends, which are correlated with but not causally related to changes in speed within LEs over time, confound the estimated broadband speed effect. Moreover, we can also use difference-in-differences techniques to reassure ourselves that, conditional on the strong controls employed, there are no within LE trends correlated with distance to the LE that could lead to spurious broadband supply effects.

We finally note that eq. (4) and eq. (5) are complementary. Adding LE x year fixed effects φ_{jt} to eq. (4) would partially absorb the identifying discontinuous variation in speed over time across

[21] It is telling that all the regulatory analysis done by Ofcom, which relies on information supplied by the broadband operators, is, indeed, conducted at the LE level, instead of at a more disaggregated level, such as street cabinets. This is because the regulator believes that the relevant market for business decisions is the LE, which is where most investments have to be sunk.

LE boundaries. Likewise, adding LE x year boundary fixed effects ψ_{kt} to eq. (5) would partially absorb the identifying continuous variation in speed over time within LEs. Because the two equations are designed to identify the broadband capitalization effect from two different types of variation, consistent estimates will be particularly indicative of their robustness.

3.1 The relationship among technology, distance and speed

As said above, we have very detailed information on the exact broadband capacity to deliver achievable speeds at a specific property at a high spatial detail, but not over the entire period. We know, however, the technology available in each LE at different points in time. We now establish the technological relationship between actual Internet speed, the technology of a LE, and the distance from a test location to the LE, using the comprehensive data set of Internet speed tests in the sub-period 2009-10. Combining both ingredients, it is possible to generate the micro-level Internet speed panel variable we require for a robust identification of the causal effect of broadband capacity on house prices.

We model broadband actual speed as a function of LE characteristics and the distance to the LE, as well as the interaction between the two. In doing so, we first need to account for a significant proportion of speed tests that are likely constrained not only by technological limitations (distance to the LE and LE characteristics), but also by the plans users have chosen to subscribe to. In other words, speed can be low not because technology is limited, but because a subscriber with small consumption choses a plan with limitations. We want to get rid of these plans so that we can unravel the true speed that a certain technology can potentially supply. To identify the plans that do not constrain broadband speed beyond the technological limitations of the LE, we run the following auxiliary regression:

$$\log(S_{ijt}) = \sum_{m=2}^{12} \alpha_m + \sum_{h=1}^{23} \alpha_h + \sum_{w=1}^{6} \alpha_w + \sum_{p=1}^{62} \alpha_p + \sum_{d=2}^{60} \alpha_d + \varphi_{jt} + \varepsilon_{ijt}, \qquad (6)$$

where S_{ijt} is the actual broadband speed test score measured at postcode i served by local exchange j at time t. α_m are month of the year effects (baseline category is January), α_h are hours of the day effects (baseline category 0h), α_w are day of the week effects (baseline category Sunday), α_p are Internet plan effects (baseline category is missing information), α_d are distance to LE effects captured by 100m bins (e.g., 2 covers distances from 150 to 250m, baseline category is 0-150m), and φ_{jt} are a set of LE-year specific fixed effects that capture unobserved LE characteristics in a given year. For the ensuing analysis, we keep observations whose α_p falls in the upper quartile, as the plans that realize the fastest actual speeds are unlikely to be constrained by the provider.

Using this sub-sample of speed tests that should be constrained only by technology, we then establish the technological relationship between available actual broadband speed S_{ijt} and distance to the relevant LE $(DIST_{ij})$ for each technological category $Q = \{\text{ADSL}, \text{ADSL} + \text{LLU}, \text{ADSL2} +\}$ in separate regressions of the following type:

$$\log(S_{ijt}) = \sum\nolimits_{m=2}^{12} \alpha_{mQ} + \sum\nolimits_{h=1}^{23} \alpha_{hQ} + \sum\nolimits_{w=1}^{6} \alpha_{wQ} + \sum\nolimits_{n=0}^{4} \alpha_{nQ} (DIST_{ij})^n + \varphi_{jQ} + \omega_{tQ} + \varepsilon_{ijtQ}. \quad (7)$$

The fourth-order polynomial is used to capture the non-linearities reported in the technical literature.[22] Since we drop 75% of the observations compared to eq. (6) and split the remaining sample into three categories in order to find technology-specific effects, we account for location and year effects separately, rather than accounting for their interaction, to save degrees of freedom in sparsely populated LEs. Based on the estimated distance decay parameters α_{nQ} and the known Q-type upgrade dates T_j^Q, it is then straightforward to predict the available actual broadband speed at any postcode *i* that is served by a LE *j* over the entire period:

$$S_{ijt} = \begin{cases} ISDN = 128 \text{ kbit/s} \ \text{ if } t < T_j^{ADSL} \\ \exp\left[\sum\nolimits_{n=0}^{4} \alpha_{nQ} (DIST_{ij})^n\right] \text{ if } T_j^Q \leq t < T_j^{Q'} . \end{cases} \quad (8)$$

This compact formulation says that, before broadband is rolled out in LE *j*, the line is served with a basic ISDN technology, as a voice telephony line is in place. Then, ADSL brings its upgraded speed at any period after T_j^{ADSL}. The decay parameters may further change if the LE additionally receives, at a certain point in time $T_j^{Q'}$, technology $Q' = \{\text{ADSL} + \text{LLU}, \text{ADSL2} +\}$.

We start by reporting the results on the physical relationship among speed, technological characteristics of the LE, and distance between the premise and the LE, as described by model (7). Our findings are shown in Table 1.

Although, due to space limitations, we do not detail the various fixed effects in the table, they all show a very reasonable behavior. The time of day is an important factor: the average connection speed reaches its peak at 5 a.m., when download speed is about 12% faster than the reference speed at midnight. It then gradually declines, with speed 3% lower at noon, 11% lower at 6 p.m. and close to 20% lower at 8 p.m., when the worst daily speed is attained. From then on, the average speed of a connection gradually increases until 5 a.m.. The day of the week also determines average speed: it is lowest over the weekend, when residential users tend to be at home. These findings are due to obvious local congestion when most people are online

[22] For a list of the factors that affect local broadband speed, see, e.g., the explanation provided by BT: http://bt.custhelp.com/app/answers/detail/a_id/7573/c/. A detailed analysis of the factors that affect the performance of ADSL networks is found in Summers (1999). We note that the choice of a fourth-order polynomial for distance was dictated by its goodness of fit. There was no gain in going towards higher orders.

simultaneously. Congestion is, thus, another facet of speed that shows striking analogies in the digital and the real worlds (see e.g. Couture et al., 2012; Duranton and Turner, 2011).

Turning to the impact of distance, which is of more direct interest for our purposes, this is shown in columns (1), (2), and (3) of Table 1 for ADSL, LLU, and ADSL 2+, respectively. Distance plays a statistically very significant role for all of them. Table 1, column (4) also runs a placebo test. The cable technology, which is available only in some parts of the country, does not rely on copper wires and does not suffer from distance-decay problems. Thus, the distance of a home from any exchange should not impact speed. Column (4) reports the results for one set of cable contracts offered by the cable provider, and, indeed, distance is found to have no impact.

One way of showing the relevance of the results is to evaluate the fit of the polynomial approximation. We estimate the distance relationships replacing the polynomial, as estimated in Table 1, with a set of 100m distance bin effects, as used in equation (6). Results are shown in Figure 3. Solid lines are the fourth-order polynomials (from Table 1) fitted into the raw data (not the dots). The dots indicate the point estimates of 100m bins obtained in separate regressions for each technology. The fit is quite striking, especially for distances up to 5 km from the LE—for greater distances, there is also more noise because there are few observations beyond that distance. We are, thus, confident that we can approximate the real speed sufficiently precisely so that attenuation bias can be ignored in equations (4) and (5). We further note that we use estimated parameters of a physical relationship that depends on distance and LE technology to approximate our speed capacity variable.

	(1)	(2)	(3)	(4)
	log of download speed (in kbit/s)			
Technology	Broadband ADSL	Broadband ADSL+LLU	Broadband ADSL2+	Cable
Distance from test postcode to LE in km	0.184	0.057	0.053	0.016
	(0.145)	(0.121)	(0.071)	(0.032)
Distance ^2	-0.293***	-0.287***	-0.491***	0.016
	(0.097)	(0.097)	(0.055)	(0.029)
Distance ^3	0.058**	0.070**	0.141***	-0.001
	(0.024)	(0.028)	(0.017)	(0.010)
Distance ^4	-0.003*	-0.005**	-0.011***	-0.001
	(0.002)	(0.002)	(0.002)	(0.001)
Constant	7.869***	8.214***	8.672***	8.334***
	(0.098)	(0.065)	(0.036)	(0.017)
LE effects	YES	YES	YES	YES
Month effects	YES	YES	YES	YES
Day of the week effects	YES	YES	YES	YES
Hour of the day effects	YES	YES	YES	YES
Year effects	YES	YES	YES	YES
r2	0.174	0.160	0.198	0.034
N	53,961	64,447	310,256	290,067

Notes: Only observations falling into the top-quartile of contracts are used in the regressions. Standard errors in parentheses are clustered on LEs. * p<0.1, ** p<0.05, *** p<0.01

Table 1: Speed results

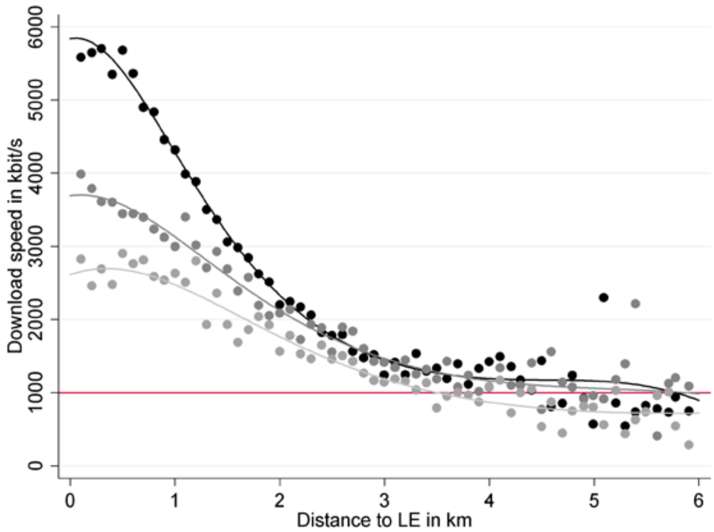

Notes: Black lines and dots indicate ADSL2+ LEs, dark (resp. light) grey lines and dots are ADSL LEs with (resp. without) LLU

Figure 3: Distance decay by LE type

These results confirm the key role played by distance. First, there is strong speed decay by distance: as a building happens to be farther away from the relevant LE, its actual speed goes down compared to another dwelling connected to the same LE with the same technology, but closer to the exchange. This phenomenon is particularly strong within 3 km (2 miles) around a LE, which is a threshold often mentioned in the technical and policy literature.[23] Second, speed decay exists for each technology, but in different ways. ADSL2+ is the newest technology (within our sample period) that can ensure the highest speeds, but it also suffers from relatively faster decay. The different sensitivity of speed to distance by technology is something that we can exploit in our main pricing models, which we discuss next.

4 Empirical findings

4.1 The impact of speed on property prices

We now give an empirical answer to our main question: Does broadband speed have an impact on property prices? Table 2 shows the result of estimating the model given by eq. (4), in columns (1-3), and by eq. (5), in columns (4-6). For both models, we first estimate the average effect of a 1 Mbit/s increase in speed, excluding (columns 1 and 4) and including (columns 2 and 5) control x year effects. We then add quadratic speed terms to allow for diminishing returns, as predicted by our theory (columns 3 and 6).

[23] See Summers (1999) and, e.g., "... like all copper technologies, the speed of ADSL2+ depends on line quality and distance; beyond 3 km from the exchange there is no real speed advantage over ordinary ADSL." http://www.worcestershire.gov.uk/cms/pdf/INCA-Beyond-Broadband.pdf.

	(1)	(2)	(3)	(4)	(5)	(6)
	log of sales price (in GBP)					
Imputed local broadband speed in Mbit/s	0.0189*** (0.0022)	0.0156*** (0.0022)	0.0254*** (0.0041)	0.0432*** (0.0018)	0.0124*** (0.0007)	0.0253*** (0.0014)
Speed^2			-0.0026*** (0.0009)			-0.0026*** (0.0002)
4^{th} order distance poly.	YES	YES	YES	YES	YES	YES
Controls	YES	-	-	YES	-	-
Control x year effects	-	YES	YES	-	YES	YES
LE effects	YES	YES	YES	-	-	-
LE boundary x year effects	YES	YES	YES	-	-	-
LE x year effects	-	-	-	YES	YES	YES
Boundary window (m)	200	200	200	-	-	-
r2	0.9485	0.9511	0.9511	0.9224	0.9317	0.9318
N	125,209	125,209	125,209	1,082,777	1,082,777	1,082,777

Notes: For columns (1-3), we identify the broadband effect from discontinuous variation in speed over time and across LE boundaries. Identification in columns (4-6) derives from a comparison of house prices to broadband supply over time and within LE areas. We further add controls on LE boundary x year effects for (1-3) and LE x year for (4-6). We present the boundary estimates for a 200m boundary window. The results for boundary windows ranging from 100m to ∞ are available in Appendix E, Table E1. Standard errors in parentheses are clustered on LE boundary x year effects in (1-3) and on LE x year cells in (4-6). * p<0.1, ** p<0.05, *** p<0.01

Table 2: Pricing results

We find positive and significant capitalization effects of broadband speed in all models. Adding control x year effects reduces the marginal speed effect from 4.3% to 1.2% when we identify from within-LE variation (columns 4 and 5). The difference is much smaller when we identify from variation across LE boundaries (1.9% vs. 1.6%; columns 1 and 2). This is the expected result because shocks to property prices are arguably less likely to be correlated with speed increases across a LE boundary within a small boundary segment (see Figure 2) than with speed increases within a LE area that depends on distance to the LE. In our preferred models (3) and (6), we find virtually identical point estimates, even though we identify from different sources of variation and samples that, in terms of observations, differ by a factor of 10. Note that we have chosen a spatial window of 200m on each side of a LE boundary in columns (1-3) as a compromise that resulted in small boundary areas that are reasonably well populated. Note, also, that we have replicated model (3) using windows of varying sizes (Appendix E1). Likewise, we have excluded varying windows from model (6) to make the samples used in (3) and (6) mutually exclusive. Because the estimates are very similar in all models, we present them in Appendix E.

Given the virtually identical point estimates in (3) and (6), we conclude that the differences in the average effects reported in columns (2) and (5) are a composition effect, as the full sample includes more properties close to LEs where the highest speeds are realized.[24] Moreover, the control x year effects seem to do a good job in capturing within-LE trends, making model (6) our

[24] When we calculate the elasticity of property prices with respect to speed, as implied by specifications (3) and (6), we obtain remarkably identical values of 0.031 at the mean of each sample.

preferred model for the counterfactual analysis, as it is estimated from our universe of property transactions and exploits the full variation in speed.

The point estimates in models (3) and (6) imply a marginal effect of 1.4% at a (post-2000) mean (real) speed of 2.2 Mbit/s. This corresponds to a 3% elasticity of property prices with respect to speed. The marginal effect of speed becomes zero at a real speed of about 5 Mbit/s, which corresponds to about 20 Mbit/s in nominal terms and roughly the 99th percentile in the overall speed distribution in our data. The implied effect on property prices at this point is 3.8% and, thus, £8,360 (≈$12,540) for a property worth £220,000 (≈$330,000, the mean house price in 2005, which is the middle point of the 2000-2010 period of Internet development we cover).[25] It is interesting to see that the marginal effect (i.e., the impact of a marginal increase in speed on net consumer surplus in our model) is about zero close to the maximum actual speed that we observe in the data. There would be no particular reason for suppliers to provide speed above the maximum observed levels in our data, as no further surplus could be created.

Using our preferred specification (6), we have produced results that show the capitalization effect by region. These are summarized in Figure 4. The left panel (in logs) shows the results as percentages, while the right panel (in levels) converts the findings in monetary rents. It is reassuring that the marginal effects in the left panel look relatively similar.[26] It seems important to acknowledge that prices differ substantially across English regions. Similar marginal capitalization effects may, therefore, imply different rents. In fact, the striking, though perhaps not surprising, result is that we get a broadband marginal monetary rent that is about twice as high in London as in any other English region. After having estimated separate effects for each region, London shows higher than average willingness to pay for broadband, but it is not an outlier in this distribution. The difference in the marginal rent is, instead, attributable to the higher house-price levels in London. Usage is probably also a lot higher in London than in the rest of the country, but competition among broadband providers is very intense too, so they

[25] This premium is comparable to, e.g., an increase in floor size of about 8 square meters, holding all other housing characteristics (e.g., the number of rooms) constant, or a reduction in distance to the nearest underground station by roughly one kilometer (Gibbons and Machin, 2005). We can compare our findings with available figures from works that have followed different approaches. Rosston et al. (2010) estimate demand from US survey data (the survey was administered online) and report that the representative household is willing to pay $48 per month for an improvement in speed from slow to very fast. Their speed variable takes only three categorical values (slow/fast/very fast), while we have the actual available speed. Still, it is reassuring to find that their consumer surplus estimates, when translated into a perpetuity using a 5% interest rate, gives $11,520 which is very close to our $12,540 estimate of the effect of going from slow to very fast. A 5% interest rate is a reasonably high discount rate as, in our data, if one buys a property into an area that has not been upgraded to the latest technology, the disadvantage compared to other areas is likely to persist over time because also in the future one would be likely to receive upgrades later. The capitalization effect thus captures an anticipated stream of rents over a relatively long period.

[26] The English regions defined in the NBS dataset are: East Anglia, East Midlands, London, North West, Northern, Out Metropolitan, Outer South East, South West, West Midlands, Yorkshire and Humberside. We do not label all of them in Figure 4 to improve readability.

cannot really price-differentiate accordingly. It is property sellers in London who ultimately receive a higher rent from broadband usage.

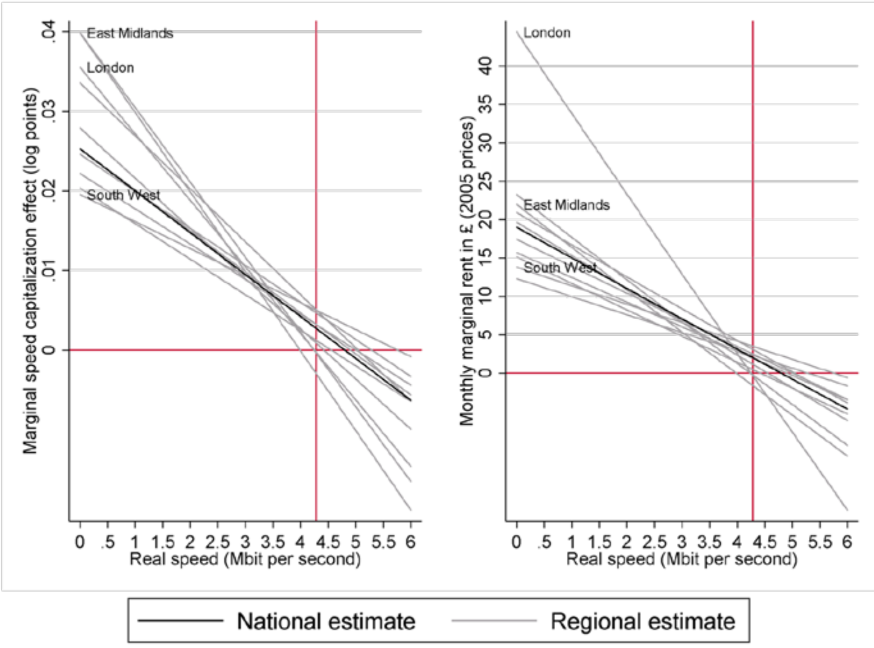

Notes: The left panel shows the marginal speed capitalization effects by regions. The right panel computes the corresponding monthly monetary rent. The monthly marginal rent R_r' is constructed as $R_r{}' = \bar{P}_r \times c/12 \times (\exp(\alpha_{r1} + 2\alpha_{r2}S) - 1)$ using the following ingredients: A 2005 adjusted mean sales price $\bar{P}_r$ in English regions recovered from the region fixed effects φ_R of an auxiliary hedonic regression of type $\log(P_{it}) = \widetilde{X}_i'\mu + \sum_{t\neq 2005}\omega_t + \varphi_R + \epsilon_{it}$; an opportunity cost of capital of c = 5%; the region-specific speed parameters α_{r1} (linear speed term) and α_{r2} (quadratic speed term) obtained from separate estimations of eq. (5) for each of the ten English regions. Grey solid lines show the respective marginal effects estimated from the regional samples. Black solid lines illustrate the marginal effect (Table 2, column 3) for the entire sample. The red vertical line indicates the 95th percentile in the (post-2000) speed distribution across the country.

Figure 4: WTP by regions

Our results do suggest that a broadband rent exists in general. Local characteristics, however, also seem to be important. The rent is rather low in regions with a higher share of low-income rural areas, which is probably where access to broadband is a problem. It seems that the benefits are relatively small where the policy maker is most likely to intervene. If the subsidies required are sufficiently low, there may still be some rationale for interventions. What also seems to be important is that the rent is declining in speed. For policy, this may imply that what is really important is to make sure that everyone gets access to some decent broadband connection. Getting access to very high speeds should, perhaps, not be the priority. This is what we analyze in the policy section. Before doing so, however, we conduct some further checks to reassure that broadband speed does, indeed, cause an increase in property prices.

4.2 Sources of identification and robustness checks

In this section we shed further light on the sources of identification that underlie the results presented above as well as their robustness.

Figure 5 illustrates the nature of the identification from discontinuous changes in speed over time across LE boundaries exploited by eq. (4). Specific to each boundary segment, we define the period before the first upgrade took place as the *BEFORE* period. The remaining period is the *AFTER* period. Within each boundary segment, we define the side of the boundary with the higher speed in the *AFTER* period as the *FAST* side (positive distance from the LE boundary). Likewise the side with the lower speed is the *SLOW* side (negative distance from the boundary). The figure pools all the raw data together, as this is the most transparent way to show the main source of variation in our identification strategy. Figure 5 shows that there was a flat distribution of speeds before the first upgrades took place (upper left). After the first upgrade, there is an evident discontinuity with higher speeds on the *FAST* sides. Note that for the purpose of illustration we keep the allocation to *FAST*/*SLOW* after the first upgrade constant over time, even thought it may change in reality. This creates some potential fuzzyness in the figure. This problem does not arise in the actual capitalization models that we estimate as we capture speed as a variable that changes continuously in space and over time. Still, Figure 5 illustrates that on average across the *AFTER* period speeds are significantly higher within the side that was first upgraded. In line with these higher speeds on the *FAST* side during the *AFTER* period, we see higher property prices on the on the *FAST* side during the *AFTER* period. There is also a notable discontinuity in the distance trend. Neither the positive difference in prices on average nor the positive discontinuity at the boundary as one moves towards the *FAST* side does exist during the *BEFORE* era. The implication is that the higher prices within the *FAST* side in the *AFTER* period are unlikely caused by time-invariant features that are specific to the *FAST* area.

We note that the endowment of various types of amenities tends to be symmetric on both sides of the boundary and there are no clear discontinuities at the boundary (see Figure C1 in Appendix C). It is therefore unlikely that the discontinuity in prices during the *AFTER* period shown in Figure 5 is caused by time trends correlated with these amenities.

There is also little evidence that our estimated effect of broadband speed on property prices is driven by changes in the composition of buyers or property characteristics. In Table C1 in Appendix C we present estimates of eq. (4) using a range of buyer or property related characteristics as dependent variables. We find no significant effect of broadband speed on whether a buyer is a first-time buyer or signs a leasehold contract, on the size of transacted properties, on whether these properties are new, have central heating or are flats (instead of houses).

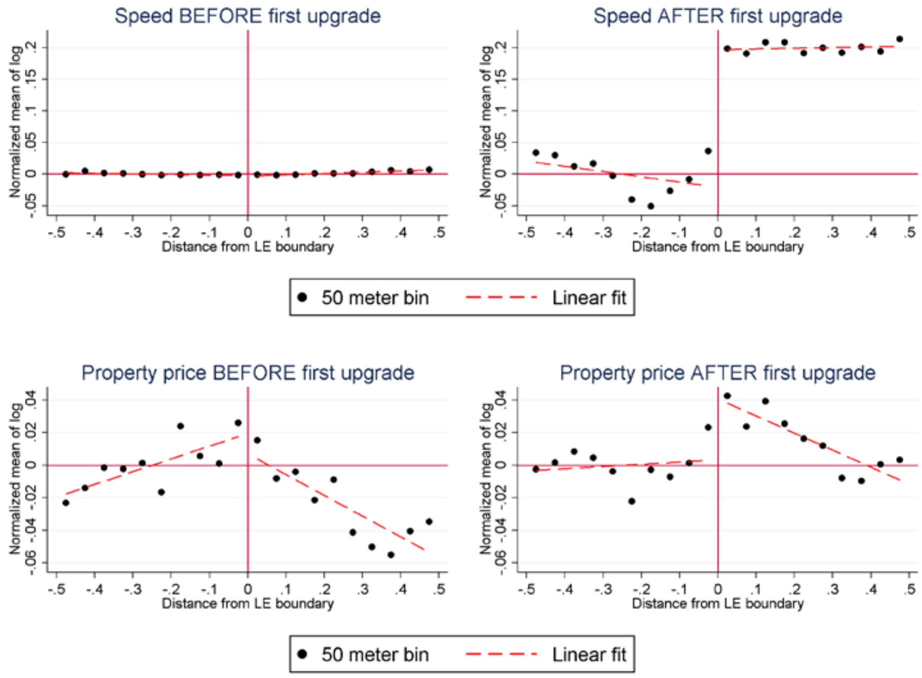

Notes: Negative distances indicate locations within the side of the boundary segment that showed lower speeds after the first upgrade of either side. Dots are means across transaction prices and real speeds within 50m distance bins.

Figure 5: Boundary discontinuities in speed in property prices

In a similar spirit, we now illustrate in a transparent way the spatiotemporal adjustment in property prices to LE upgrades within LE areas in Figure 6. Our methodology is explained in detail in Appendix D, where we discuss a reduced-form difference-in-differences (DD) specification, expanded to account for spatial heterogeneity and for a temporal structure in the treatment effect of a LE upgrade. Figure 6 allows us to investigate how the relationship between property prices and distance to the LE changes up to three years prior to the ADSL upgrade (*PRE* placebo effects) as well as up to three years after the ADSL upgrade (*POST* treatment effects), in each case relative to the period three or more years before the upgrade. We note that Figure 6 shows the average effect across all ADSL upgrades estimated conditional on LE and year effects. All estimated *PRE*-treatment ADSL effects are near to zero and most are even slightly negative. Property prices did not tend to be higher close to LEs before the ADSL upgrade, despite notable correlations between various forms of amenities and LE distance (see Figure D1 in Appendix D). With the upgrade, prices increase close to the LEs, which is in line with a significant positive effect of real broadband speed that declines in distance from the LE. While there is a slight orientation over the three years preceding the ADSL activation towards a more negative distance gradient, the level shift *after* the upgrade is very substantial. The effects for the three *POST* periods are very consistent, and it seems fair to conclude that these cannot be explained by trends that existed *prior* to the upgrade.

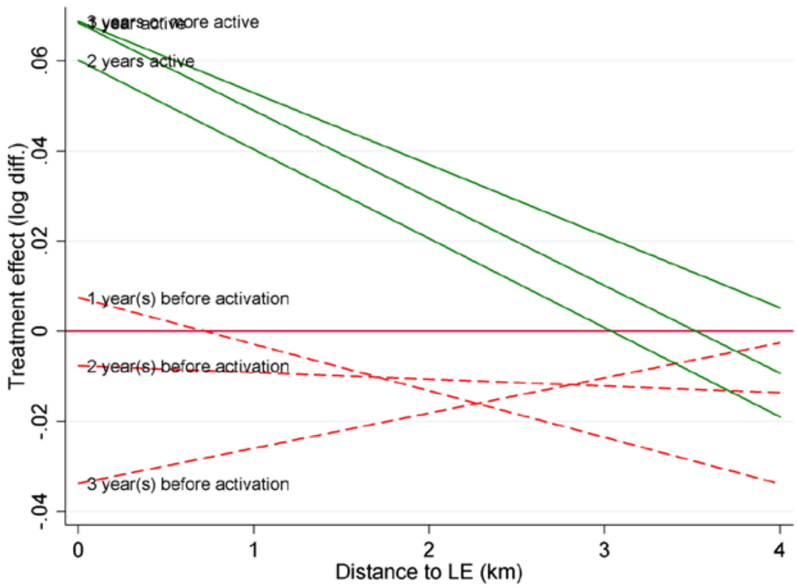

Notes: Red dashed (green solid) lines show difference-in-differences estimates for periods before (after) the ADSL upgrade took place.

Figure 6: Difference-in-differences results with spatiotemporal variation: ADSL

We now return to the empirical models of Section 4.1. To support our benchmark model results and to substantiate our economic interpretations of the findings, we have run a series of additional models. The results are summarized in Table 3. To control for a long-run trend correlated with distance to the LE and not absorbed by control x year effects, we add an interaction between the fourth-order distance to LE variables and a linear time trend in column (1). This is a strong control as it is likely to partially absorb the effect of speed upgrades if capitalization occurs smoothly over time. In line with the discontinuous pattern in Figure 6, the speed effect, however, remains remarkably close to the benchmark model, pointing to speed capitalization effects that occur discontinuously in time.

Our results could be biased in the presence of externalities at a very disaggregated level, for instance at the building level. One possibility is that speed might attract particular people to a block of flats first, and subsequent buyers might be enticed by the proximity to those original buyers rather than by speed per se. To reduce this possibility, we rerun our model excluding flats, thus concentrating only on detached, semi-detached or terraced houses where only a single family could move. Results in column (2) of Table 3 are virtually identical to those reported in column (6) of Table 2 (similarly for the model with boundary discontinuities, not reported here for the sake of brevity).

Because LLU and ADSL2+ are both advancements that started only in 2005, it is possible to divide our sample to identify the speed effect from variation that stems from two separate technological innovations. A priori, results could go either way. Prior to 2005, email and browsing were the prevalent Internet activities for residential users, while phenomena such as

YouTube or Facebook were only limited. The older applications were, however, much less bandwidth intensive, in a period when available bandwidth was also much more restricted. While broadband speed is clearly very important today (because of changes in complementary technology), actually, at the margin, the willingness to pay for additional Mbit/s could be either higher or lower in the early days compared to more recent periods, as supply was much more constrained by technology. Column (3) of Table 3 uses transactions up to 2004, when most ADSL activations occurred. Likewise, column (4) uses transaction from 2005 onwards and, thus, exploits LLU and ADSL2+ activations. Results are very much in line with our benchmark model, as the differences between periods are not marked enough to be a source of alarm.

	(1)	(2)	(3)	(4)	(5)	(6)
	log of sales price (in GBP)					
	OLS	OLS	OLS	OLS	OLS	2SLS
Imputed local broadband	0.0269***	0.0255***	0.0273***	0.0214***	0.0316***	0.0288***
speed in Mbit/s	(0.0014)	(0.0014)	(0.004)	(0.0062)	(0.0021)	(0.0015)
Speed^2	-0.0018***	-0.0027***	-0.0023*	-0.0014**	-0.0038***	-0.0036***
	(0.0003)	(0.0003)	(0.0013)	(0.0007)	(0.0003)	(0.0003)
4th order distance poly.	YES	YES	YES	YES	YES	YES
4th ord. pol. x (year – 2000)	YES	-	-	-	-	-
2 Mbit/s pre-ASDL cap	-	-	-	-	YES	-
LE x year effects	YES	YES	YES	YES	YES	YES
Controls x year	YES	YES	YES	YES	YES	YES
Period	1995-10	1995-10	1995-04	2005-10	1995-10	1995-10
Property type	All	Houses	All	All	All	All
r2	0.933	0.935	0.91	0.89	0.932	0.932
N	1,082,777	932,878	729,133	353,644	1,082,777	1,082,777

Notes: In column (1), we add an interaction between the fourth-order distance to LE variables and a linear time trend to account for within-LE trends in property prices that are accidently correlated with distance to the LE. In column (2) we exclude flats. In column (3), we identify the simple ADSL speed upgrade effects in the earlier period (up to 2004), and in column (4) the combined effects from LLU and ADSL2+ upgrades (after 2005). In column (5), we use a different speed panel variable that accounts for the 2 Mbit/s cap for the period prior to 2006. In column (6), we use three indicator variables for ADSL/LLU/ADSL2+ as well as three interactions of these indicator variables with LE distance as predictors for Speed and Speed^2. Standard errors in parentheses clustered on LE x year cells. * p<0.1, ** p<0.05, *** p<0.01.

Table 3: Robustness checks

Because we have no access to speed-test data from before 2008, we are not able to fully control for some technological improvements that occurred to the basic ADSL technology. In its early years, ADSL speed was capped at 2 Mbit/s, and this constraint was removed only in 2006, allowing for the maximum nominal speed of 8 Mbit/s. Our best possible attempt to approximate the respective technological parameters is to estimate equation (7) using speed tests of users who subscribed to plans that cap the maximum speed at 2 Mbit/s. In column (5), we assign values implied by this speed-distance function to all transactions that occurred after ADSL activation, but before 2006 or LLU. The results are qualitatively identical and quantitatively similar to those of our benchmark model.

One could argue that our estimated engineering relationship between speed, technology, and distance is rather sophisticated (though it is sufficient that one has access to a website that performs the test without knowing the underlying formula, or that the available speed is known to the local estate agent that then transmits the information to prospective buyers). Also, we rely on speed tests that are initiated by users. To address these concerns, we use ADSL/LLU/ADSL2+ indicator variables, plus the interactions with LE distance as predictors for *Speed* and *Speed^2* in a 2SLS model in column (6). This way, we restrict the identifying variation to stem purely from LE technology and distance, which is a fairly transparent structure for identification that we also use in the difference-in-differences models (see Figure 6 and Appendix D). In Table E3 in Appendix E we apply the same strategy to the boundary specification and also consider a spline distance approach to predict real speed. All results consistently show that our findings do not depend on the functional form derived on the engineering analysis in section 3.1, although the latter is our preferred specification as it does not depend on ad hoc assumptions and generates precise estimates.

4.3 Heterogeneity and capitalization channels

In this section we engage with some ancillary predictions of our conceptual model. To allow for heterogeneity in the willingness-to-pay for the Internet (modelled via a_j in our conceptual model in Section 2.3) we augment eq. (5) as follows:

$$\log(P_{ijt}) = \sum_{m=1}^{2} \alpha_m (S_{ijt})^m + \sum_{m=1}^{2} \left((S_{ijt})^m \times A_i\right) \beta_m^A + \sum_{n=1}^{4} \tau_n (DIST_{ij})^n + X_i' \mu_t + \varphi_{jt} + \epsilon_{ijt}, \quad (9)$$

where A_i is a vector of time-invariant characteristics of property i, capturing population characteristics (average income), urbanization (share of urban land, labor market accessibility), amenities (school quality, proximity to rail stations, restaurant density, retail density) and internet services (Amazon evening delivery, Uber, number of retailers dispatching online orders), which are discussed in more detail in Appendix B. $\beta_{m=1,2}^A$ are the respective parameter vectors capturing spatial heterogeneity in the WTP for speed. Because the estimates of these interaction terms between the quadratic speed term and the relatively highly correlated locational variables are difficult to interpret, we relegate the presentation and a more detailed discussion of the results to Table F1 in Appendix F. Briefly summarized, we find that the marginal effect of real broadband speed is larger in more urban areas with higher incomes and more amenities, suggesting that such areas are inhabited by households with a relatively high willingness-to-pay for speed. In this context it is worth noting that in Table F2 in Appendix F we rerun our preferred models (columns 3 and 6 in Table 2) separately for buyers who did and did not purchase a property for the first time. We do not find any notable difference between the two groups on the willingness to pay for speed, both in the overall sample, and at the boundaries.

We also find evidence for a complementarity between broadband speed and availability of local internet services. In Figure 7 we illustrate the distribution of the estimated marginal speed effects across properties distinguishing between locations where local internet services such as Amazon evening delivery or Uber are supplied (solid lines) and those where they are not (dashed lines). Evidently, the distributions are shifted to the right in areas with such services, suggesting that higher broadband speed is valued more by buyers in those areas. In Figure F2 in Appendix F we similarly show that the marginal speed effect increases in the number of grocery chains that dispatch online orders to a certain location. Even though such spatial heterogeneity may be partially attributable to differences in socioeconomic status, these results represent a significant addition to the scarce evidence on how the value of broadband depends on the supply of complementary internet services (Forman et al., 2008).

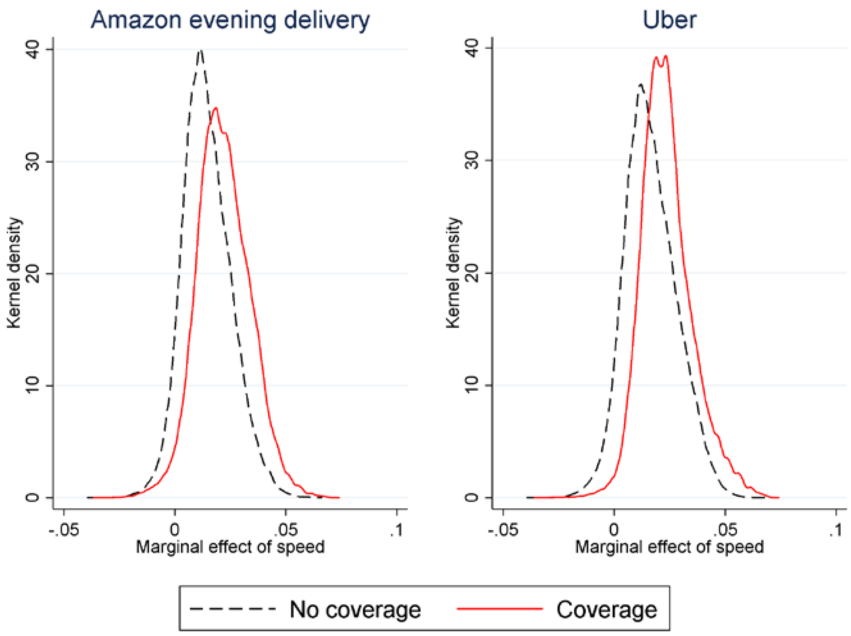

Note: The marginal speed effect is defined as: $\frac{\partial logP_{ijt}}{\partial S_{ijt}} = \alpha_1 + A_i\beta_1^A + 2\alpha_2 S_{ijt} + 2(S_{ijt} \times A_i)\beta_2^A$ and derived from model (3) in Table F1 in Appendix F, which allows the speed effect to vary in population characteristics (average income), urbanization (share of urban land, labor market accessibility), amenities (school quality, proximity to rail stations, restaurant density, retail density) and available internet services (Amazon evening delivery, Uber, number of retailers dispatching online orders). Kernel is Epanechnikov. A Kolmogorov-Smirnov test rejects the null of the distribution for “coverage” and “no coverage” to be the same (KS = 0.267 for Amazon and KS = 0.268 for Uber; p-value < 0.01 for both).

Figure 7: Marginal speed effects in areas with and without local internet services

The result that having fast and reliable internet is more valuable where delivery of online orders is fast and more retailers can bring groceries at home seems to suggest that the consumer surplus we are measuring arises from consumption of internet services at home (as opposed to amenities such as internet cafes). This is supported by the magnitude of our estimated consumer surplus, which is in line with studies that have focused on broadband consumption at home using different methods (see footnote 25). To substantiate this interpretation we provide a direct test of the ancillary prediction of our model that faster broadband should not only lead to positive capitalization effects, but also to higher penetration rates. In particular, penetration, defined as the ratio of the number of households connected to broadband over all households in a certain area, should increase in broadband speed at a decreasing rate (see eq. (3)).

	(1)		(2)	
	Penetration (share)			
	ADSL		Cable	
Imputed local broadband speed in Mbit/s	0.0779***	(0.0066)	0.0028	(0.0018)
Speed^2	-0.0111***	(0.001)	-0.0005	(0.0003)
LE effects	YES		YES	
TTWA x year effects	YES		YES	
LE trend effects	YES		YES	
Cable coverage	ALL		>65%	
r2	0.354		0.53	
N	70,074		13,228	

Notes: Penetration rate is defined as the ratio of the number of households connected to broadband over all households in a certain area. Study period is 2005-2010. To accommodate LE trends we estimate the model in first differences including LE effects. Standard errors in parentheses clustered on LEs. * p<0.1, ** p<0.05, *** p<0.01.

Table 4: Penetration results

In Table 4 we use a strongly balanced panel of penetration rates available quarterly across LEs, ranging from the last quarter of 2005 to the second quarter of 2010, the same period as used in model (4) of Table 3. Because we cannot exploit within-LE variation, we cannot add LE x year effects to control for unobserved macroeconomic shocks at the LE level. Still, to strengthen identification, we allow for TTWA x year effects and individual LE trends (on top of LE effects).[27] As the model predicts, we find a positive speed effect on penetration that diminishes in speed (column 1). To evaluate whether unobserved shocks (e.g., gentrification) that impact broadband demand (penetration) and upgrade decisions (and, thus, speed) are driving the results, we also conduct a falsification test using cable broadband penetration rates as the dependent variable. Cable is a completely separate technology that should not, per se, be affected by the speed of the ADSL-based network. As cable is available only in some parts of the country, we restrict the analysis to those LEs with high potential cable coverage according to the Ofcom definition (more than 65% of households in a given catchment area are "passed" by cable and, thus, have potential access to cable). Reassuringly, we do not find a significant effect of speed in this placebo test (column 2). Because unobserved macroeconomic shocks that are correlated with our speed measure and increase broadband demand should also show up in higher cable

[27] Travel to Work Areas (TTWAs) are self-contained labor market areas defined by the Office for National Statistics. At least 75% of an area's resident workforce work in the area and at least 75% of the people who work in the area also live in the area. According to the 2007 definition there are 243 TTWAs in the UK.

penetration rates, we conclude that the ADSL penetration effect is unlikely to be spurious. These results support our main finding that households value broadband. Moreover, they suggest that the benefits from broadband are at least partially incurred through consumption of broadband at home, and not only through the attraction of amenities such as internet cafes, or places of cultural production and consumption that depend on a decent broadband connection to operate.

5 Evaluation of the EU Digital Agenda

In this section, we propose an evaluation of the EU Digital Agenda. As discussed in the Introduction, by 2020, every EU household should have access to at least 30 Mbit/s. In order to conduct the counterfactuals, we use the estimated capitalization effects from the hedonic regressions in order to make welfare comparisons.

The conclusion about willingness to pay for broadband upgrades requires us to think carefully about the nature of heterogeneity in broadband demand. As put by Kuminoff et al. (2013, p. 1038) it is legitimate to make welfare comparisons using results from hedonic regressions only when the analyst can reasonably answer 'yes' to the following questions: "Do the data describe a single geographic market connected by a common hedonic price function? Was the gradient of the price function constant over the duration of the study period? Are the "treated" houses in the sample representative of the population of interest?" As for the single geographic market, we have already shown how to extend our estimates to make them specific to local markets. As for the time-variation of the gradient of the price function, we did not find any particularly worrying variation at least between the pre- and post-2005 periods that we could test in Table 3. The final point is instead more controversial and harder to tackle in a reduced-form framework like ours. For sure, the buildings in our sample seem to be representative of the population. Figure 1 already gave some information about this, and we run several other reassuring tests in this direction.[28] From our tests for speed effects on buyer and property characteristics (see Table C1 in Appendix C) we also know that buyers and properties before and after speed upgrades are similar. However, people moving into properties may sort themselves according to their preference for broadband speed and, depending on whether fast internet connections are abundant or scarce, the recovered willingness to pay by marginal buyers may under- or overstate the average willingness to pay. The virtually immediate capitalization of increases in broadband available speed (see Figure 6) seems to suggest that fast connections during our study period were relatively scarce, thus, the sorting effect will likely be upward.

[28] We find that our sample of property transactions closely resembles the full population of postcodes in terms of the kernel distribution of distances to the nearest LE, which is the most important determinant of speed.

As already discussed, we do not find any notable difference on the willingness for speed between first-time buyers and other buyers (see Table F2 in Appendix F). While this is reassuring with respect to sorting, we are aware of the limitations of our data in that we lack additional household characteristics. Keeping this limitation in mind, we now offer guidance on how to interpret our results. In our policy experiment, we are going to increase Internet speed available locally to some households. If a household was interested in this higher level of speed, but could not find it as it was not available for various reasons (for instance, because of the high fixed costs to deploy a faster technology in that area), we can indeed use our results to estimate the benefit to that household from a speed increase. However, if a household was not interested in the Internet, and decided not to subscribe, it is also likely that this household will be reluctant to subscribe also when we change the broadband speed. This is particularly relevant as the EU target states that *every* household should have at least 30 Mbit/s, and thus broadband supply would have to be expanded considerably. Using the results from existing subscribers to inform the welfare attributable to these households is likely to lead to an overestimation of the true benefits from speed. For these reasons, we propose below to distinguish between benefits from "speed upgrades" and those from "coverage upgrades". This distinction keeps the welfare results separate between households with and without a broadband connection, as the former results are probably more credible than the latter.

We now present our policy experiment. In order to provide an estimate of the costs and benefits of the proposed targets, one would need to first establish the counterfactual—that is, what speeds will be reached by 2020 *without* interventions? The targets themselves must be interpreted, as the EU guidelines are not very clear. For instance, "having access" may simply mean that the target speed is technologically available in a certain area or, alternatively, that each household must effectively subscribe to that target speed.

In order to move forward, we have to make some explicit assumptions. We propose the following methodology. First, we take advantage of a useful and timely report published in November 2013 by the DCMS, the UK government's department responsible for the Internet. The report forecasts the distribution, by density decile, of the broadband speeds that will be reached in England by 2020 in the absence of interventions. This is shown in Table 5.

Density decile	1	**2**	3	4	5	6	7	8	9	10
Speed (Mbit/s)	3.88	**32.23**	75.84	120.06	169.18	218.40	250.41	277.96	294.88	332.77

Source: DCMS (2013)

Table 5: Predicted broadband speeds in England by 2020

We make some small adjustments to account for the fact that the DCMS refers to the sum of upload and download speeds, while the EU Digital Agenda refers only to download speeds.[29] It turns out that, with a very good degree of approximation, the EU target implies bringing every household to at least the average speed of the second decile of the speed distribution. We use this information to anchor our data. Of course, the broadband market will evolve between now and 2020. Our maintained hypothesis is, however, that the current relative distribution of speeds is informative as to where the market will go. Someone currently in the bottom decile of the distribution will also be at the bottom of the distribution in 2020, and so forth. Everyone will likely move towards higher speeds, but in a proportional manner.

If one is prepared to accept our assumption, then the rest of the exercise follows quite naturally. Since we can estimate benefits from broadband at the LE level, we take the 2010 distribution of speeds in England at the same LE level (see Appendix G for more details). Within this distribution, we take the average speed of the second decile, which becomes our "2010 target-equivalent" speed to which everybody should aspire by 2020, according to the Digital Agenda. We thus interpret the policy "as if" everybody should have access to at least the speed of the second decile, which we denote as S_{DA}.

Having identified the "2010 target-equivalent" speed in our data, we turn to the benefits for each LE, as this is where the targets might have an impact. To calculate LE-specific estimates of the broadband benefits, we use eq. (9). For the counterfactual exercise we allow the effect of speed to vary in income I (calculated at the 2005 ward level) and urbanization U (share of urbanized area within a 1 km^2 grid) (see column 1 in Table F1 in Appendix F). We considered models with richer sets of interactions (columns 2 and 3 in the same table), but because most amenities are highly correlated with income and urbanization these models produced implausible outliers in the speed effects for various LEs without adding much explanatory power.

We can calculate LE-specific estimates of the broadband benefits as:

$$\alpha_{1j} = \alpha_1 + \beta_1^I I_j + \beta_1^U U_j,$$

$$\alpha_{2j} = \alpha_2 + \beta_2^I I_j + \beta_2^U U_j,$$

where I_j and U_j are the means of the properties transacted within LE j and β_1^I, β_1^U, β_2^I and β_2^U are part of the vectors β_1 and β_2 in eq. (9) moderating the interactions between speed and speed^2 and the locational characteristics. The marginal effect is:

[29] See European Parliament (2013). The upload speed is roughly 10% of the download speed.

$$\left(\frac{\partial logP}{\partial S}\right)_j = \alpha_{1j} + 2\alpha_{2j}\, S.$$

To get to the marginal rent, we require some LE-level mean prices that account for differences in income and urbanization. One approach would be to use local means estimated in a similar way to the regional prices used in Figure 4 (see Figure notes for details), just at a more local level (using finer fixed effects). The other approach is to make the price income and urbanization specific—i.e., estimate prices as function of U and I:

$$\log(P_{jt}) = a_I I_j + a_U U_j + \widetilde{\mathrm{X}}_i'\mu + \sum\nolimits_{t\neq 2005} \omega_t + \epsilon_{jt}. \qquad (10)$$

The advantage of this approach is that it is possible to express the rent entirely as a function of S, U, I.[30]

For each local exchange, we also know the average speed (S) and the proportion of households (x) that have access to broadband. In every LE, we proceed as follows:

- If $S > S_{\mathrm{DA}}$, then no speed upgrade is needed in that LE. If one interprets "access" in the Digital Agenda as "technological availability," then nothing should happen in that LE. If, instead, one interprets the target more strictly—i.e., literally all households should actually subscribe to broadband with a minimum speed—then the unconnected households will need to be covered as long as x is less than 100% in that LE. For these households, the benefit is calculated by giving them the target speed S_{DA} (starting from a basic connection, corresponding to ISDN, as they will have a telephone line): we call this possible benefit "coverage upgrade."
- If $S < S_{\mathrm{DA}}$ in a given LE, the households with broadband will need a speed upgrade, leading to an increase in benefits corresponding to an increase in speed from S to S_{DA} in that LE: we call this benefit "speed upgrade." As above, if the unconnected households also must be connected, the "coverage upgrade" benefit is similarly calculated by giving them the target speed S_{DA} (starting from a basic connection).

Having described our methodology to get an estimate of the benefits from the upgrade, we need to have a view about the corresponding costs. We borrow this information from existing studies. While there are many technologies that could achieve very high speeds, it is agreed that fiber has the most promising chances of being rolled out to the mass market (and has already started in some places across England). According to how deeply fiber is deployed, the most expensive solution is fiber to the home (FTTH). A slightly less expensive solution that could still allow for

[30]It is $r = \frac{c}{12} \times \exp(a_0 + a_U U + a_I I) \times (\exp(\alpha_1 + 2\alpha_2 S + \beta_1 I + 2\beta_2 S \times I + \gamma_1 U + 2\gamma_2 S \times U) - 1)$.

very high speeds is fiber to the building (FTTB). The cost of rolling out these technologies varies by area, as they are typically cheaper in densely populated areas and more expensive in rural areas. The European Investment Bank (EIB) gives an estimate of the average NPV cost, per technology and per area, in the EU.[31] These are reported in the top two rows of Table 6.

The results of the benefits for S_{DA} are shown in the third and fourth rows of Table 6. The results by LE are aggregated by area type, to make them directly comparable with the cost estimates. We present the findings distinguishing between the gains predicted for those who already have broadband, and will just need an "upgrade" to close the speed gap, as opposed to the gains accruing to those that currently do not have broadband but will need to be "covered" to meet the target. This corresponds also to two different interpretations of the EU digital agenda.

We believe this is the most transparent way to organize and discuss our findings. Benefits are calculated as an average per household in each LE. Although we do account for differences in urbanization and income among LEs, we cannot control for other sources of unobserved heterogeneity. Hence, the "upgrade" results are probably the more credible, as they refer to households that are interested in broadband and already subscribe to it. These results are also in line with the looser interpretation of the targets, whereby technology must be available, but subscription decisions are left to individuals.

	Population density in residents/km^2		
Costs/*Benefits* per HH (GBP)	> 500 (Urban)	> 100 & < 500 (Suburban)	< 100 (Rural)
Cost (FTTH)	416	1,018	2,522
Cost (FTTB)	310	885	2,301
Speed upgrade benefit	*668*	*337*	*393*
Coverage upgrade benefit	*8,815*	*4,690*	*3,145*
LEs affected	183	257	1,075
Households affected:			
Upgrade ($S < S_{DA}$)	851,880	387,743	584,874
Coverage (x < 100%)	5,066,954	432,781	319,468

Notes: Cost estimates by density categories are taken from the EIB (Hätönen, 2011).

Table 6: Estimated costs and benefits for the 30 Mbit/s Target of the EU Digital Agenda

The "coverage" results apply, instead, to households that currently do not have a basic version of broadband, even in areas where fast broadband is available. This could be due to affordability issues, in which case our results on coverage would stand if appropriate subsidies were also given to those households. But one could also argue that these households are simply not

[31] The cost assessment is based on a combination of population densities, technology and labor costs. It refers to the fixed costs per household needed to bring a technology to a certain area. We use the 2010 average EUR/GBP exchange rate to calculate the figures for England. See Hätönen (2011) and Gruber et al. (2014) for more details on the approach. Notice that, should mobile technology be used to bring high-speed broadband to rural areas, instead of fiber, this would affect only the cost rows in Table 6, not the estimated benefits which are related to speed only, not to the delivering technology.

interested in broadband, and never will be, unless additional actions are also taken—e.g., to increase their degree of digital literacy (especially for households with older people). If one takes a stricter interpretation of the Digital Agenda, such that *every* household must have broadband of a certain minimum speed, one cannot just ignore the issue. Instead of arguing one way or another, we give each set of results separately.

Households in urban areas clearly pass the cost-benefit test. The benefits of the upgrade per household are already sufficient to cover its cost, even with the most expensive FTTH technology. As for suburban households, FTTB might be considered, but the benefits of the speed upgrade alone are still less than 40% of its cost. If a small percentage of the coverage benefits could also be realized, one could also argue for FTTB in suburban areas. Rural areas are, instead, the most problematic: this is where costs are highest and benefits lowest. The benefits from the speed upgrade are about 15% of the cost of bringing fast broadband. Only if one is willing to accept that at least two thirds of the coverage benefits will also be realized, then the case for FTTB passes a cost-benefit test under the stricter interpretation of the Digital Agenda in rural areas.

The last rows in Table 6 give some sense of the total impact of the policy. Almost two hundred LEs would need to be upgraded in urban areas, but they would affect large numbers of households, as the population density is high. Overall, the speed upgrade would affect just over 1.8m households, and possibly fewer than 1.3m if rural areas were thought to fail the cost/benefit test. Connecting the unconnected is, instead, a more ambitious goal, which puts the number of affected households well over 5m. These large differences are due to the ambiguity in interpreting the policy targets.

Our welfare assessment is based on the costs to supply broadband—and net household benefits from using it—over and above the price paid to Internet Service Providers. We have been silent so far on the actual broadband price that subscribers pay. This is not a problem if the price is competitive, so that ISPs themselves make no extra rents. If, though, there were private rents to ISPs, then our analysis would underestimate welfare effects since ISPs' profits are excluded from our study.

We finish this exercise by commenting on the possible direction of bias in our results. First, our whole approach depends on estimating broadband value from property scarcity prices. If the property market were oversupplied instead, then we would systematically underestimate consumer surplus from broadband consumption, as sellers would not be able to capture broadband rents. In this respect, it is well documented that the supply of properties in England is severely constrained by the planning system (Hilber and Vermeulen, 2015). More land is covered by greenbelts that prevent expansion of developed areas (and in some areas even by golf courses) than by housing. This restriction of developable land leads to the economically paradoxical combination of skyrocketing house prices (more than tripled in England and more than quadrupled in London over the past 15 years) and historically low construction levels (Cheshire, 2014). Still, it is safe to say that our estimates should provide a lower bound to net consumer surplus.

Second, and more relevant for the policy exercise, the relative scarcity of properties may be lower in rural areas compared to urban areas. If that were the case, then the underestimation would be more severe for the former than for the latter. While it is beyond the scope of the current work to use a measure of the tightness of the property market, we have information about the number of days it takes, on average, to sell a property from when it is first put on sale, which is an indication of how many active prospective buyers there are for that property. On the basis of this imperfect metric, there is no evidence that the supply of properties in rural areas is considerably more elastic than in urban areas.[32]

Third, if buyers anticipated broadband speed increases over time, the present value of a technological upgrade would be reduced, and we would similarly underestimate the consumer surplus. We find in our data that the sequence in which LEs where upgraded to ADSL, LLU, and ADSL2+ was similar, implying that relative speed advantages should tend to persist over time. Also, when we run DDs for each technological upgrade (see Figure 6 and Appendix C2), we find some genuine discontinuities in property prices associated with the various generations of broadband technologies, which reveals that the benefits of the introduction of ADSL and of its subsequent upgrades were not fully anticipated by consumers.

Fourth, we calculate the benefits from the digital targets in a certain LE by eventually changing only the speed in that LE, and keeping all other parameters constant. While this is not particularly controversial for urbanization, we also keep income constant. If, say, broadband became available in rural area A, and some rich people were induced, as a consequence, to move to that area A from some other area B, we would have to use their income to evaluate the policy (starting with the speed level available in their original area B). Since none of this information is available, our policy experiment is valid to the extent that there is very low mobility among LEs.

Fifth, we estimate only the private gains from residential broadband Internet. Therefore, we may be missing various positive network externalities linked to high-speed communications. It is notable, however, that urban areas already pass the cost-benefit test and rural areas fail by a large margin. Because most economic activity concentrates in urban areas, it is unlikely that the

[32] For instance, in January 2007, before the financial crisis, it took, on average, 86 days to sell a property in Greater London, the most densely populated area in England, and 95 days to sell one in rural Devon. After the crisis, these went up to 178 days and 206 days, respectively, but the relative ratio did not change (see “Time on the market report for England”, http://www.home.co.uk/guides/).

qualitative conclusions from our policy exercise would change if, for instance, the effects on firms were taken into account.

Sixth, and as we acknowledged more generally at the beginning of this section, we cannot tell what part of our property capitalization effects could be due to pure sorting. This is why we decided to be as transparent as possible by presenting the benefit results split into two parts. Perhaps the results are less credible at the extensive margin (bringing people to fast Internet for the first time) than at the intensive margin (giving a faster connection to those who already use the Internet). If this is the case, as already argued above, our most convincing estimates of broadband benefits are those capturing the speed upgrade, while the coverage upgrade estimates should be taken with more caution.

6 Conclusions

This paper evaluates the extent to which broadband speed is capitalized into house prices. We estimate consumer surplus associated with broadband Internet speed by using microdata on property prices in England between 1995 and 2010. We find a 3% elasticity of property prices with respect to speed at the mean of the speed distribution in our data. Because of significant diminishing returns to speed, this elasticity applies only to marginal changes and properties with average Internet connections. Upgrading a property from a normal (8 Mbit/s) to a fast (24 Mbit/s) connection increases the value, on average, by 1%. This is still a large effect. We argue that this is a good measure of net consumer surplus associated with broadband usage. This is true as long as properties are scarce and sellers are, thus, able to extract buyers' consumer surplus, or else our results would underestimate the impact on consumer surplus. We also find considerable heterogeneity of these benefits in each area where the Internet is locally deployed. We then use the estimates to evaluate the benefits associated with government initiatives to upgrade digital speed. We show that urban areas pass a cost-benefit test of current EU policy proposals, while the case for these policy interventions is not very strong in rural areas.

Since it is largely urban areas that pass a cost-benefit test, the question arises: Why do ISPs supply sub-optimal speed in those areas, where there seems to be a willingness to pay that is in excess of costs? The reason is that the broadband rent goes to the "wrong" economic agent. The broadband speed rent is, in fact, appropriated by the seller, not by the ISPs. The ISPs supply broadband according to supply and demand conditions in the broadband market, which is largely a competitive one. But these conditions do not necessarily reflect the scarcity rents that exist in the property market. To upgrade their local networks, ISPs need to recover substantial fixed costs (especially for fiber) over the relevant catchment area. ISPs can recover these fixed costs only in part via the premium prices charged to subscribers, since they are still restrained by the competitive landscape.

An implication of our results is that there may be a coordination problem among sellers and landlords in the undersupplied areas that pass the cost-benefit tests, perhaps because they are unaware or, most likely, because of their fragmentation. While it would be collectively rational for these sellers and landlords to get together and pay some of the ISPs' delivery costs of upgraded technologies—as, then, their properties would become more valuable—free-riding problems make this scenario unlikely. As with other infrastructures, the coordination problem, therefore, rationalizes the public delivery of broadband to undersupplied areas in combination with levies charged to sellers and landlords to recover part of the costs. The political economy of the housing-markets literature suggests that homeowners and landlords would support such initiatives as long as the anticipated capitalization gain exceeds the infrastructure levy (Ahlfeldt et al., 2014; Dehring et al., 2008; Fischel, 2001; Oates, 1969).

Bibliography

Ahlfeldt, G. M., and Kavetsos, G. (2014). Form or Function? The impact of new sports stadia on property prices in London. *Journal of the Royal Statistical Society A, 177*(1), 169-190.

Ahlfeldt, G. M., Moeller, K., Waights, S., and Wendland, N. (2014). Game of zones: The political economy of conservation areas. *CESifo Working Paper No. 4755*.

Baum-Snow, N. (2007). Did Highways Cause Suburbanization? *The Quarterly Journal of Economics, 122*(2), 775-805.

Baum-Snow, N., and Kahn, M. E. (2000). The effects of new public projects to expand urban rail transit. *Journal of Public Economics, 77*(2), 241-263.

Chay, K. Y., and Greenstone, M. (2005). Does Air Quality Matter? Evidence from the Housing Market. *Journal of Political Economy, 113*(2), 376-424.

Chen, Y., and Savage, S. J. (2010). The Effects of Competition on the Price for Cable Modem Internet Access. *Review of Economics and Statistics, 93*(1), 201-217.

Cheshire, P. C. (2014). Turning houses into gold: the failure of British planning. *CentrePiece, Spring 2014*, 15-18.

Couture, V., Duranton, G., and Turner, M. A. (2012). Speed. *Working Paper*.

Czernich, N., Falck, O., Kretschmer, T., and Woessmann, L. (2011). Broadband Infrastructure and Economic Growth. *Economic Journal, 121*(552), 505-532.

Davis, L. W. (2004). The Effect of Health Risk on Housing Values: Evidence from a Cancer Cluster. *American Economic Review, 94*(5), 1693-1704.

DCMS. (2013). *UK Broadband Impact Study. Department for Culture, Media and Sport, November 2013*. London.

Dehring, C. A., Depken, C. A., and Ward, M. R. (2008). A direct test of the homevoter hypothesis. *Journal of Urban Economics, 64*(1), 155-170.

Duranton, G., Morrow, P. M., and Turner, M. A. (2014). Roads and Trade: Evidence from the US. *The Review of Economic Studies, 81*(2), 681-727.

Duranton, G., and Turner, M. A. (2011). The Fundamental Law of Road Congestion: Evidence from US Cities. *American Economic Review, 101*(6), 2616-2652.

European Parliament. (2013). *Entertainment x.0 to Boost Broadband Deployment. PE 507.479, Directorate-General for Internal Policies, October 2013*. Brussels.

Faber, B. (2014). Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System. *The Review of Economic Studies, 81*(3), 1046-1070.
Figlio, D. N., and Lucas, M. E. (2004). What's in a Grade? School Report Cards and the Housing Market. *American Economic Review, 94*(3), 591-604.
Fischel, F. A. (2001). *The Homevoter Hypothesis* (Vol. Cambridge, MA.): Harvard University Press.
Forman, C., Ghose, A., and Goldfarb, A. (2008). Competition Between Local and Electronic Markets: How the Benefit of Buying Online Depends on Where You Live. *Management Science, 55*(1), 47-57.
Forman, C., Goldfarb, A., and Greenstein, S. (2012). The Internet and Local Wages: A Puzzle. *American Economic Review, 102*(1), 556-575.
Gibbons, S., and Machin, S. (2005). Valuing rail access using transport innovations. *Journal of Urban Economics, 57*(1), 148-169.
Gibbons, S., Machin, S., and Silva, O. (2013). Valuing school quality using boundary discontinuities. *Journal of Urban Economics, 75*(0), 15-28.
Goolsbee, A., and Klenow, P. J. (2006). Valuing Consumer Products by the Time Spent Using Them: An Application to the Internet. *American Economic Review, 96*(2), 108-113.
Greenstein, S., and McDevitt, R. C. (2011). The broadband bonus: Estimating broadband Internet's economic value. *Telecommunications Policy, 35*(7), 617-632.
Greenstone, M., and Gallagher, J. (2008). Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program. *The Quarterly Journal of Economics, 123*(3), 951-1003.
Gruber, H., Hätönen, J., and Koutroumpis, P. (2014). Broadband access in the EU: An assessment of future economic benefits. *Telecommunications Policy, 38*(11), 1046-1058.
Hätönen, J. (2011). The economic impact of fixed and mobile high-speed networks. *EIB Papers, 16*(2), 30-59.
Hilber, C. A. L., and Vermeulen, W. (2015). The impact of supply constraints on house prices in England. *Economic Journal, onlinefirst*.
Kolko, J. (2012). Broadband and local growth. *Journal of Urban Economics, 71*(1), 100-113.
Kuminoff, N. V., Smith, V. K., and Timmins, C. (2013). The New Economics of Equilibrium Sorting and Policy Evaluation Using Housing Markets. *Journal of Economic Literature, 51*(4), 1007-1062.
Lewis, M. (2014). *Flash Boys*. New York: W.W. Norton.
Linden, L., and Rockoff, J. E. (2008). Estimates of the Impact of Crime Risk on Property Values from Megan's Laws. *American Economic Review, 98*(3), 1103-1127.
Nardotto, M., Valletti, T., and Verboven, F. (2015). Unbundling the Incumbent: Evidence from UK Broadband. *Journal of the European Economic Association, 13*(2), 330-362.
Nevo, A., Turner, J., and Williams, J. (2015). Usage-Based Pricing and Demand for Residential Broadband. *Econometrica, forthcoming*.
Oates, W. E. (1969). The Effects of Property Taxes and Local Public Spending on Property Values: An Empirical Study of Tax Capitalization and the Tiebout Hypothesis. *Journal of Political Economy, 77*(6), 957-971.
Summers, C. K. (1999). *ADSL: standards, implementation, and architecture* (Vol. 4). Boca Raton, Florida: Crc Press.

Appendices

Appendix A: Evolution of broadband in England

Figure A1 below shows the evolution of the availability of ADSL (first panel), LLU (second panel), and ADSL2+ (last panel) in every area of England over the study period. The red dots show the location of all the LEs in England. The first panel shows that ADSL became ubiquitous by the end of the period, though upgrades happened at different points in time in different areas. The second and last panels show that LLU and ADSL2+ did not diffuse everywhere, and a considerable part of the country (the hatched areas, which are concentrated in the rural parts of the country) did not attract sufficient economic interest from providers to bring faster broadband there.

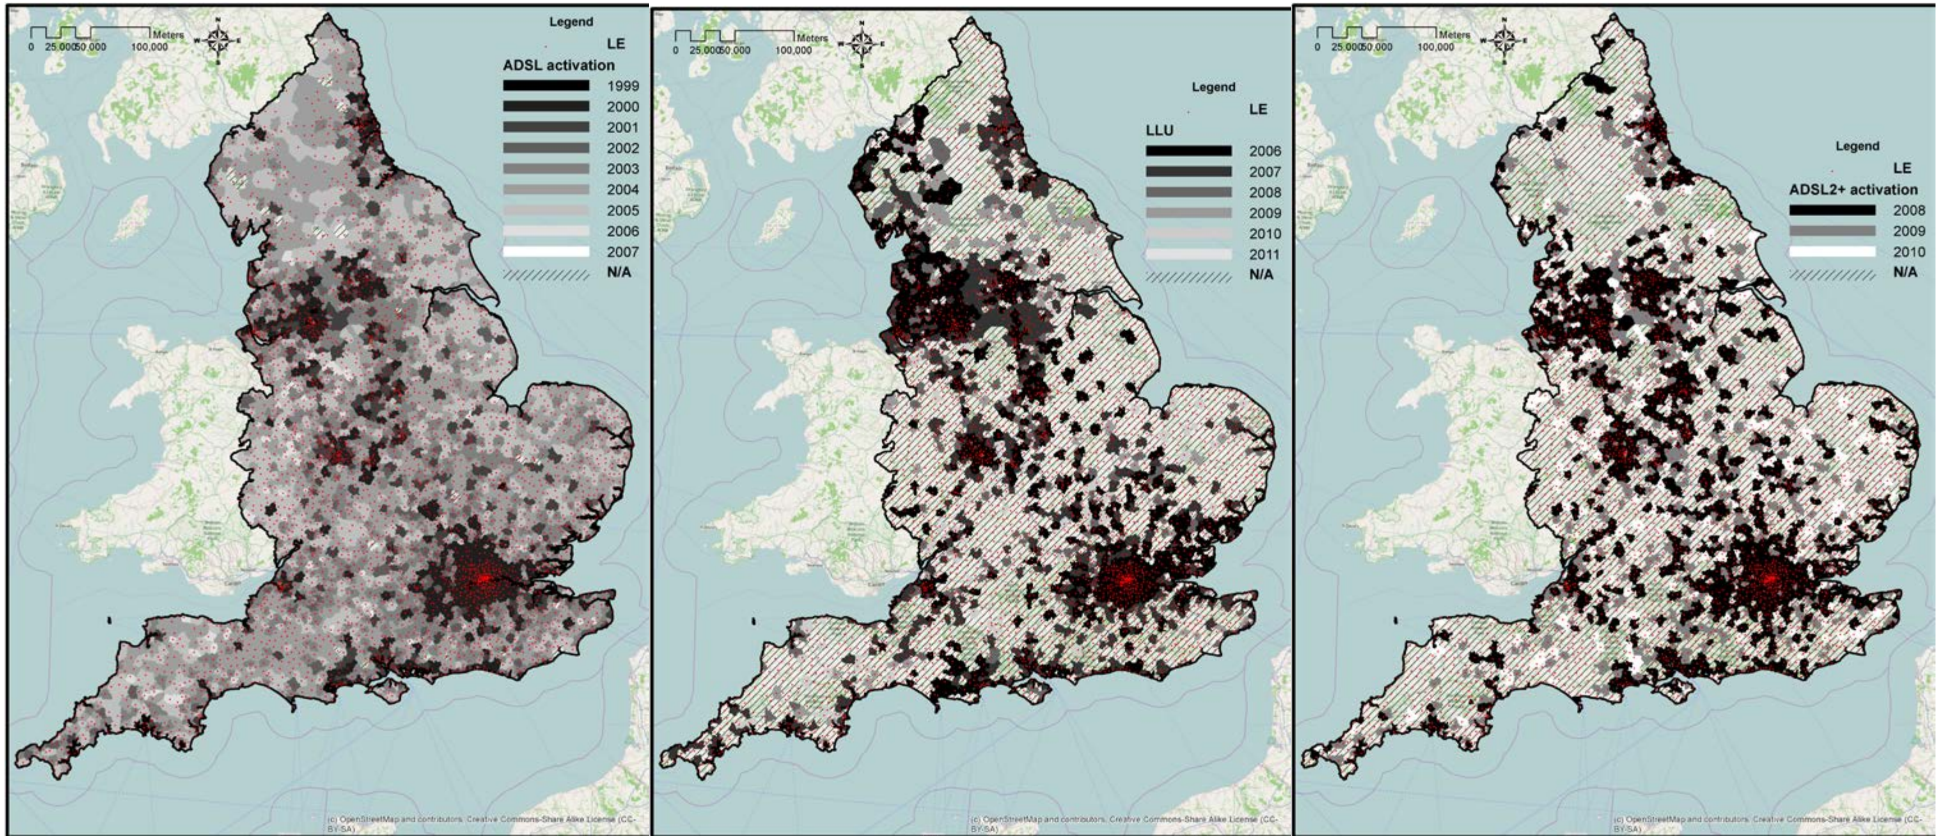

Note: Red dots illustrated the location of LEs

Figure A1: The evolution of ADSL, LLU and ADSL2+ in England (1999-2010)

Appendix B: Data description

In this appendix, we introduce the additional non-broadband speed-related covariates we use in the capitalization regressions in more detail. Table B1 provides summary statistics of the most important variables in our dataset and Table B2 a summary description of the full list of variables employed. See, also, Ahlfeldt et al. (2014).

Neighborhood characteristics

The main variables used for estimating capitalization effects of neighborhood characteristics are median income and ethnic composition. The income data provide a model-based estimate of median household income produced by Experian for Super Output Areas of the lower level (LSOA). This is assigned to the transaction data based on postcode. The data on ethnicity were made available by the 2001 UK Census at the level of Output Area (OA). Shares of each of the 16 ethnic groups and a Herfindahl index[33] were computed to capture the ethnic composition of neighborhoods.

Environmental variables

The environmental variables capture the amenity value of areas e.g. designated as natural parks, various features of the natural environment, and different types of land cover and use.

Geographical data (in the form of ESRI shapefiles) for UK National Parks, Areas of Outstanding Natural Beauty, and National Nature Reserves are available from Natural England. National Parks and Areas of Outstanding Natural Beauty are protected areas of countryside designated because of their significant landscape value. National Nature Reserves are "established to protect sensitive features and to provide 'outdoor laboratories' for research." Straight-line distances to these designations were computed for the housing units as geographically located by their postcodes. Furthermore, density measures that take into account both the distance to and the size of the features were created. We apply a kernel density measure (Silverman, 1986) with a radius of 2km, which is considered to be the maximum distance people are willing to walk (Gibbons and Machin, 2005).

The location of lakes, rivers and coastline is available from the GB Ordinance Survey. The distance to these features is also computed for the housing units from the transaction data. The UK Land Cover Map produced by the Centre for Ecology and Hydrology describes land coverage by 26 categories, as identified by satellite images. We follow Mourato et al. (2010), who construct nine broad land cover types from the 26 categories. Shares of each of these nine

[33] The Herfindahl index (HI) is calculated according to the following relation: $HI = \sum_{i=1}^{N} s_i^2$, where s_i is the share of ethnicity i in the LSOA, and N is the total number of ethnicities.

categories in 1km grid squares are calculated, and the housing units take on the value of the grid square in which they reside.

The generalized Land Use Database (GLUD) available from the Department for Communities and Local Government gives area shares of nine different types of land use within Super Output Areas, lower level (LSOA). These nine types are domestic buildings, non-domestic buildings, roads, paths, rail, domestic gardens, green space, water, and other land use. These shares are assigned to the housing units based on the LSOA in which they are located.

Amenities

The locational amenities variables capture the benefits a location offers in terms of accessibility, employment opportunities, school quality, and the proximity of cultural and entertainment establishments.

Employment accessibility is captured both by the distance to Travel to Work Area (TTWA) centroid and by a measure of employment potentiality. TTWAs represent employment zones, and the distance to the center of these zones is a proxy for accessibility to employment locations. A more complex measure of accessibility is the employment potentiality index. This is computed at the Super Output Area, lower level (LSOA) and represents an average of employment in neighboring LSOAs, weighted by their distance.[34]

Key Stage 2 (ages 7–11) assessment scores are available from the Department for Education at the Super Output Area, middle layer (MSOA). School quality is captured at the house level by computing a distance-weighted average of the KS2 scores of nearby MSOA centroids.

Geographical data on the locations of motorways, roads, airports, rail stations and rail tracks are available from the GB Ordinance Survey. Distances were computed from housing units to motorways, A-roads, B-roads and rail stations to capture accessibility. Buffer zones were created around the motorways and roads along with distance calculations to rail tracks and airports in order to capture the unpleasant noise effects of transport infrastructure.

Further data on local amenities were taken from the Ordinance Survey (police stations, places of worship, hospitals, leisure centers), GeoLytix (retailers) and OpenStreetMap (cafés, restaurants/fast food outlets, museums, nightclubs, bars/pubs, theaters/cinemas, kindergardens and monuments, attractions). The number of listed buildings was provided by English Heritage. Kernel densities for these amenities were computed for housing units using a

[34] The potential is defined as $\sum_M E_M \exp(-\tau d_{iM})$, where E_M is the total employment at MSOA M, d_{iM} is the distance between property i and MSOA M, and τ is a decay parameter chosen to reflect the spatial decay in bilateral commuting probabilities as discussed in more detail by Ahlfeldt et al. (2014).

kernel radius of 2km and a quadratic kernel function (Silverman, 1986). The radius of 2km is consistent with amenities having a significant effect on property prices only when they are within walking distance.

These data are complemented by variables capturing the eligibility for Amazon evening deliveries and Uber fleet services. The variables were constructed in GIS based on maps (Uber) and postcode sector mappings (Amazon) available at the respective company websites. As a further measure of online services density we approximate the number of grocery chains[35] within a 10-mile ride, the usual distance over which retailers dispatch online orders.[36]

Variable	N	Mean	S.D.	Min	Max
Log price (£)	1,082,777	11.653	0.659	8.780	15.021
Real speed (Mbit/s)	1,082,777	1.454	1.389	0.128	5.844
Distance from LE (km)	1,082,777	1.430	0.938	0.002	14.414
Distance from LE boundary (km)	1,082,777	0.937	0.674	0.000	5.873
ADSL (at LE)	1,082,777	0.545	0.498	0.000	1.000
LLU (at LE)	1,082,777	0.249	0.432	0.000	1.000
ADSL2+ (at LE)	1,082,777	0.097	0.296	0.000	1.000
Household income (£/year)	1,082,777	27,436	8,304	3,952	85,646
Share urbanized land (1km^2 grid)	1,082,777	0.585	0.276	0.000	1.000
Employment potential	1,082,777	128,835	179,156	357	1,406,765
Distance from rail station (km)	1,082,777	2.501	2.782	0.003	59.894
Key-stage 2 test score (MSOA)	1,082,777	27.351	0.980	15.233	31.098
Restaurant density	1,082,777	0.858	2.275	0.000	99.762
Retail density	1,082,777	0.581	0.603	0	8.49507
Amazon evening delivery	1,082,777	0.441	0.496	0	1
Retailers delivering online orders	1,082,777	6.625	1.491	0	8
Uber	1,082,777	0.178	0.383	0	1

Table B1: Descriptive statistics of key variables

[35] The considered chains include Aldi, Asda, Booths, Budgens, Co-op, Costco, Lidl, Marks and Spencer, Morrisons, Sainsbury's, Tesco, Waitrose and Wholefoods.

[36] We approximate the 10-mile road network distance by $\frac{10}{\sqrt{2}}\sim 7.07$ miles straight-line distance (Ballou et al., 2002).

Dependent Variable	
Price	Log transaction price in GBP of a property from the Nationwide Building Society (NBS).
Independent Variables	
Housing information	Set of property variables from the NBS including: Number of bedrooms, number of bathrooms, floor size (in square meter), new property (dummy), building age (years), tenure (leasehold/freehold), central heating (full: gas, electric, oil, solid fuel), central heating (partial: gas, electric, oil, solid fuel), garage (single or double), parking space, property type (detached, semi-detached, terraced, bungalow, flat-maisonette).
Neighborhood information	Set of neighborhood variables including: median income (2005, LSOA level), share of white population at total population (2001 census, output area level), share of mixed population at total population (2001 census, output area level), share of black population at total population (2001 census, output area level), share of Asian population at total population (2001 census, output area level), share of Chinese population at total population (2001 census, output area level), Herfindahl index of ethnic segregation (including population shares of White British, White Irish, White others, Mixed Caribbean, Mixed Asian, Mixed Black, Mixed other, Asian Indian, Asian Pakistani, Asian others, Black Caribbean, Black African, Black other, Chinese, Chinese other population, 2001 census output area).
Environment Characteristics and Amenities	Set of locational variables processed in GIS including: National Parks (distance to, density), Areas of Outstanding Beauty (distance to, density), Natural Nature Reserves (distance to, density), distance to nearest lake, distance to nearest river, distance to nearest coastline, land in 1km square: Marine and coastal margins; freshwater, wetland and flood plains; mountains, moors and heathland; semi-natural grassland; enclosed farmland; coniferous woodland; broad-leaved/mixed woodland; urban; inland bare ground.
Other amenities	Set of locational variables created in GIS including: Average key stage 2 test score (MSOA averages as well as interpolated in GIS), distance to electricity transmission lines, A-Roads (distance to, buffer dummy variables within 170m), B-Roads (distance to, buffer dummy variable within 85m), motorway (distance to, buffer dummy variable within 315m; buffer distances refer to the distance were noise of maximum speed drops down to 50 decibel), distance to all railway stations, distance to London Underground stations, distance to railway tracks, distance to bus stations, distance to airports, densities of cafés, restaurants/fast food places, museums, nightclubs, bars/pubs, theaters/cinemas, grocery stores, kindergardens, monuments (memorial, monument, castles, attraction, artwork), hospitals, sports/leisure centers, police stations and worship locations, distance to Travel to Work Areas, employment potentiality.
Online services	Amazon evening delivery (0,1 dummy), Uber (0,1 dummy), the number of retailers delivering online orders.

Table B2: Variable description

Appendix C: BDD analysis

In this appendix we complement the main paper by illustrating how outcomes other than property prices vary spatially across LE boundaries and spatiotemporally across LE boundaries and time. We also provide additional detail on how we generate Figure 5.

1. Other amenities at the boundary

A popular validation exercise in the boundary discontinuity design (BDD) literature is to test for discontinuities in alternative spatial variables that potentially determine the outcome measure but are not related to the phenomenon of interest (e.g. Gibbons, et al., 2013). As we identify from variation that is discontinuous in space *and* time, thus being able to control for time-invariant characteristics, we are primarily interested in whether or not discontinuities in observable amenities exist across LE boundaries to sense how likely it is that unobserved trends correlated with these amenities confound our estimates.

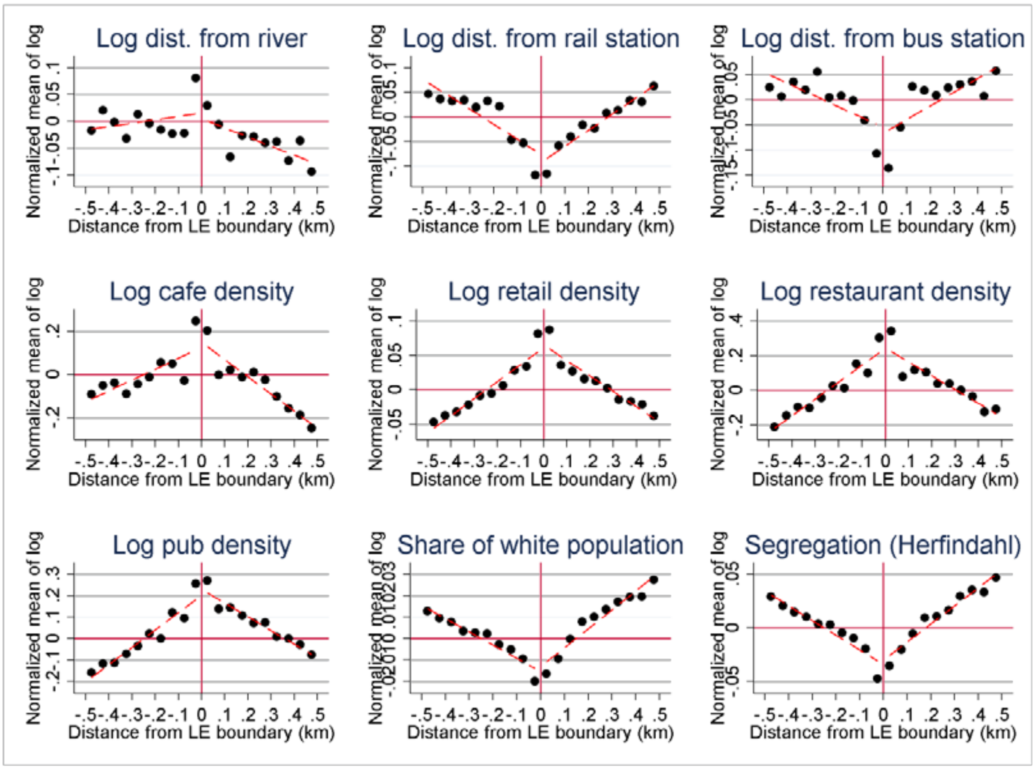

Notes: Negative distances indicate locations within the side of the boundary segment that showed lower speeds after the first upgrade of either side. Dots are means across observations within 50m distance bins. All variables are normalized to have a zero mean within the slow side (negative distances). Distances are straight-line distance to the nearest feature. Density measures use Gaussian kernels with a radius of 2km and a bandwidth selected according to Silverman (1986).

Figure C1: Boundary effects in geographic features

Figure C1 examines how a selection of time-invariant amenities, covering a mix of natural amenities, consumption amenies and neighborhood characteristics, change on either side of a LE boundary (the running variable Distance is explained in the figure notes and in more detail in section C2). While we find that locations close to the LE boundaries differ from other locations, the endowment is usally symmetric on both sides of the boundary, and there are no clear discontinuities at the boundary. It is therefore unlikely that the discontinuity in prices that we identify in the BDD is caused by time trends correlated with these amenities. We note that in the capitalization models we allow for arbitrary changes in the implicit prices of these (and other) amenities over time.

2. Spatiotemporal discontinuities: Before-after comparison of discontinuities

We now explain how we create Figure 5 in the main text. For each boundary pair, we identify the first date when an upgrate on either side of the boundary leads to a speed differential between the two sides of the boundary. The speed is uniform before the first upgrade (during the narrowband ISDN era; the upgrade is usually if not always to ADSL). Specific to each boundary segment, we define the period before this first upgrade took place as the BEFORE period. The remaining period is the AFTER period. Within each boudnary segment *k*, we define the side of the boundary with the higher speed in the AFTER period as the FAST side. Likewise the side with the lower speed is the SLOW side. We define distance from the boundary as the running variable. Within the SLOW side, the running variable takes negative values. Likewise, it takes positive values within the FAST side. The running variable thus describes a continuous move from the SLOW side to the FAST side, crossing the boundary at a 0 value within each boundary segment *k*. We group properties into 50m bins defined based on this running variable, i.e. -500 to -450, … -50 to 0, 0 to 50, … 450 to 500 m. For each bin, we compute the mean across the distribution of BEFORE speeds at properties within this bin (upper left in Figure 5). For each bin, we compute the mean across the distribution of AFTER speeds at properties within this bin (upper right in Figure 5). For each bin, we compute the mean across the distribution of transaction prices of properties sold during the BEFORE period (bottom left in Figure 5). For each bin, we compute the mean across the distribution of transaction prices of properties sold during the AFTER period (bottom right in Figure 5). We normalize (log) speeds and (log) prices to have zero means within the SLOW side. We find a significant discontinuities in speeds and prices only during the AFTER period, which is in line with a significant speed effect.

3. Spatiotemporal discontinuities: Alternative outcomes

As a falsification test that corresponds to the standard tests for boundary discontinuities in cross-sectional BDDs we are interested in whether other outcome variables systematically adjust where and when speed increases due to LE upgrades. For this purpose, we present estimates of eq. (4) using various alternative dependent variables instead of log of property price in Table C1. The models are otherwise comparable to model (3) in Table 2. We do not find a significant speed effect on any of the considered alternative outcomes, making it more likely that our baseline model captures a genuine effect of speed on property prices.

	(1) First time buyer (1=yes)	(2) Leasehold (1=yes)	(3) Log Number of bedrooms	(4) New Property (1=yes)	(5) Central heating (1=yes)	(6) Flat or maisonette (1=yes)
Imputed local broadband speed in Mbit/s	-0.0015	0.0032	-0.0037	-0.0027	0.0003	-0.0021
	(0.0098)	(0.0039)	(0.0048)	(0.0032)	(0.0057)	(0.0036)
Speed^2	0.0022	-0.0010	0.0013	0.0010	-0.0002	0.0008
	(0.0021)	(0.0009)	(0.0012)	(0.0009)	(0.0012)	(0.0008)
4^{th} order distance poly.	YES	YES	YES	YES	YES	YES
Control x year effects	YES	YES	YES	YES	YES	YES
LE effects	YES	YES	YES	YES	YES	YES
LE boundary x year eff.	YES	YES	YES	YES	YES	YES
Window	200	200	200	200	200	200
r2	0.3652	0.8579	0.7490	0.7056	0.3051	0.8717
N	125,209	125,209	125,209	125,209	125,209	125,209

Notes: The dependent variable is excluded from the control x year effects. Except for the dependent variable, models are otherwise identical to Table 2, column (3). Standard errors in parentheses are clustered on LE boundary x year effects. * p<0.1, ** p<0.05, *** p<0.01

Table C1: Alternative dependent variables

Appendix D: DD analysis

Given the impact of distance on broadband speed, an upgrade of a LE can be viewed as an event that should exert spatially variant effects on nearby property prices. The effect of this event on property prices can, thus, be analyzed using quasi-experimental research designs that have become popular in the literature. In this appendix, we complement the empirical analysis presented in the main paper using a reduced-form empirical specification that is a mix of hedonic modeling, panel econometrics, and a DD method, which accommodates multiple treatment dates and spatial heterogeneity in the treatment effect within an area served by a LE. We first show that other observable amenities tend to be systematically correlated with distance from the LE, highlighting the need to exploit variation over time to identify a speed effect from within-LE variation. Then we present the empirical framework that we use for the full sample where identification arises from within-LE variation. Finally we discuss the empirical findings.

1. Spatial distribution of amenities around the LE

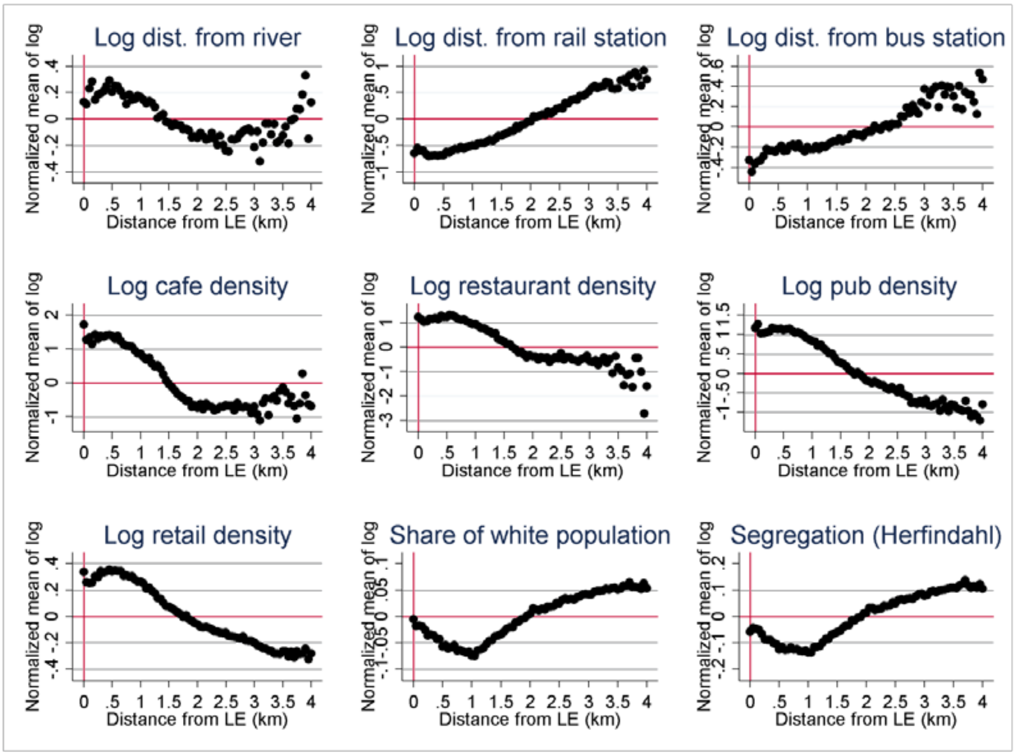

Notes: All variables rescaled to have a mean of zero. 4 km is approximately the 99th percentile in the distribution of observations across distance from the LE. Each black dot represents the mean within a 50 m distance bin. Density measures use Gaussian kernels with a radius of 2km and a bandwidth selected according to Silverman (1986).

Figure D1: Spatial distribution of amenities around the LE

Figure D1 illustrates the correlation between the distance from the LE and various amenities covering natural amenities, consumption amenies and neighborhood characteristics. In keeping with intuition, LEs tend to be located away from rivers and close to transport hubs and consumption amenities. These pronounced correlations highlight the need to exploit variation over time to control for unobserved amenities when establishing the relationship between speed and property price. Another insight from Figure D1 is that there may be property price trends over time that are correlated with LE distance, because LE distance is correlated with various amenities. We therefore allow implicit prices of all property characteristics (including the amenities reported here) to vary arbitrarily over time in our empirical models.

2. Empirical framework

The point of departure of our DD analysis summarized in Figure 6 is the following specification:

$$\log(P_{ijt}) = \sum\nolimits_{Q} \beta_{0Q}(POST_{jt}^{Q}) + \sum\nolimits_{Q} \beta_{1Q}(POST_{jt}^{Q} \times DIST_{ij}) + \beta_3 DIST_{ij} + \mathrm{X}_i'\mu + \varphi_j + \omega_t + \varepsilon_{ij}$$

where P is the sales price of a property that sells in postcode i served by LE j in year t, X_i' is a vector of structural, location and neighborhood variables and μ is a vector of implicit hedonic prices. φ_j is a fixed effect for whether a property is located within the catchment area of a LE j, ω_t is a year fixed effect and ε_{ijt} a random error term. $POST_{jt}^{Q}$ are 0,1 indicator variables indexing whether at time t, LE j had been upgraded to quality level $Q = \{\text{ADSL}, \text{LLU}, \text{ADSL2}+\}$. The treatment effect of a certain type of LE upgrade Q on property prices at a given distance from an upgraded LE is given by $\beta_{0Q} + \beta_{1Q}DIST_{ij}$. The DD comparison relative to LEs that were not upgraded and the period before the upgrade is, thus, made at every distance from the upgraded LEs.

A typical concern in DD analyses are temporal trends that are correlated with but not causally related to the treatment. Identification, in general, cannot be considered credible if changes in property prices near to LEs following an upgrade can be explained by (relative) trends in the neighborhoods that existed prior to the upgrade. The concern is relevant in our case because the assignment of the LE upgrade is not technologically random. Therefore, we expand the spatial DD model to allow for a temporal structure in the treatment effect of a LE upgrade.

In the first step, we allow for additional spatially varying DD effects for each of the three years immediately preceding an upgrade. Because we do not expect capitalization effects in anticipation of an upgrade, these effects can be viewed as placebo-treatment effects. We estimate the following model:

$$\log(P_{ijt}) = \sum\nolimits_{Q} \beta_{0Q} POST_{jt}^{Q} + \sum\nolimits_{Q} \beta_{1Q} \left(POST_{jt}^{Q} \times DIST_{ij}\right) + \sum\nolimits_{Z} \sum\nolimits_{Q} \beta_{0ZQ} \left(PRE_{Zjt}^{Q}\right) + \sum\nolimits_{Z} \sum\nolimits_{Q} \beta_{1ZQ} \left(PRE_{Zjt}^{Q} \times DIST_{ij}\right) + \beta_3 DIST_{ij} + X_i' \mu + \varphi_j + \omega_t + \varepsilon_{ijt},$$

where PRE_{Zjt}^{Q} indexes a LE x year cell Z years before a Q-type upgrade of LE j. Note that these *PRE* effects provide a DD comparison relative to LEs that where not upgraded and the period four or more years before an activation. In a further expansion, we replace the *POST* effects with separate DD effects for each of the two first years subsequent to an upgrade and a residual category that contains all subsequent years.

3. Empirical results

We report in Table D1 the results on spatiotemporal trends around the upgrade dates. To keep the tabular presentation compact, we report parametric results for models in which we add the *PRE* placebo DD effects, but no separate *POST* effects (column 2 in Table D1). In the graphical illustration for the ADSL upgrade in Figure 6 in the main text, we also allow DD effects to vary by years following designation. To save space, we do not show the corresponding figures for the LLU and ADSL2+ upgrades, but we discuss the results next.

The pattern of time-varying LLU effects is as follows. All *POST* effects show the expected pattern with a positive level shift that flattens out towards the fringe of the LE. The effect increases notably from the first to the second *POST* period and moderately afterwards. Two of the three *PRE*-effects are not in line with a successful falsification test at first glance. The effects are positive, and one shows a notable negative slope. A closer inspection reveals, however, that the *PRE* effects decline towards the activation date. Also, the negative slope tends to disappear over time. Pre-trends, thus, are negatively correlated with the treatment and are reversed just at the time of the upgrade, which makes a particularly strong case for impact.

The ADSL2+ effects show a similar pattern. In the model with separate *PRE*-effects (where the comparison is made relative to four and more years before activation), the ADSL2+ *POST* effect turns out to be negative at all distances to the LE. This is not the expected result, even though there is negative decay, as expected. The *POST* effect is, however, significantly more positive than any of the three *PRE* effects, in all areas that are relatively close to the LE. Moreover, the earlier *PRE* effects show a positive distance trend, which is reversed only one year before the ADSL2+ activation. As with the LLU effects, the inspection indicates that pre-trends are negatively correlated with the treatment, which strengthens the sense of impact.

	(1) log of sales price (in GBP)		(2) log of sales price (in GBP)	
ADSL active	0.082***	(0.005)	0.071***	(0.004)
LLU active	0.026***	(0.004)	0.056***	(0.005)
ADSL2+ active	0.010***	(0.003)	-0.004	(0.004)
ADSL x DIST	-0.018***	(0.002)	-0.018***	(0.002)
LLU x DIST	-0.006***	(0.002)	-0.008***	(0.003)
ADSL2+ x *DIST*	-0.003**	(0.001)	-0.001	(0.002)
PRE_1^{ADSL}			0.012***	(0.003)
PRE_2^{ADSL}			-0.004	(0.003)
PRE_3^{ADSL}			-0.030***	(0.003)
PRE_1^{ADSL} x DIST			-0.011***	(0.002)
PRE_2^{ADSL} x DIST			-0.002	(0.002)
PRE_3^{ADSL} x DIST			0.007***	(0.002)
PRE_1^{LLU}			0.036***	(0.006)
PRE_2^{LLU}			0.041***	(0.005)
PRE_3^{LLU}			0.048***	(0.004)
PRE_1^{LLU} x DIST			-0.006	(0.004)
PRE_2^{LLU} x DIST			0.000	(0.002)
PRE_3^{LLU} x DIST			-0.005**	(0.002)
PRE_1^{ADSL2+}			-0.012**	(0.006)
PRE_2^{ADSL2+}			-0.018***	(0.003)
PRE_3^{ADSL2+}			-0.020***	(0.003)
PRE_1^{ADSL2+} x DIST			-0.001	(0.004)
PRE_2^{ADSL2+} x DIST			0.004**	(0.002)
PRE_3^{ADSL2+} x DIST			0.007***	(0.002)
LE Effects	YES		YES	
Year Effects	YES		YES	
Controls	YES		YES	
Distance to LE	YES		YES	
r2	0.916		0.916	
N	1,070,197		1,070,197	

Notes: Standard errors in parentheses are clustered on LEs. * p<0.1, ** p<0.05, *** p<0.01.

Table D1: Difference-in-differences with spatial variation

Appendix E: Robustness checks

In this section we present the results of robustness checks of our baseline empirical models not reported in the main paper for brevity. In particular we evaluate the robustness of the results to varying the sample size with respect to distance from LE boundaries and different functional forms imposed on the speed – LE distance relationship.

1. Varying distance from boundary windows

Table E1 below presents estimates of eq. (4) for varying boundary window sizes. The models are otherwise identical to model (3) in Table 2. The results remain close to the baseline model in all specifications.

	(1)	(2)	(3)	(4)	(5)
	log of sales price (in GBP)				
Imputed local broadband	0.0266***	0.0264***	0.0264***	0.0254***	0.0228***
speed in Mbit/s	(0.0012)	(0.0016)	(0.0023)	(0.0041)	(0.0069)
Speed^2	-0.0030***	-0.0031***	-0.0030***	-0.0026***	-0.0024*
	(0.0002)	(0.0003)	(0.0005)	(0.0009)	(0.0016)
4^{th} order distance poly.	YES	YES	YES	YES	YES
Control x year effects	YES	YES	YES	YES	YES
LE effects	YES	YES	YES	YES	YES
LE boundary x year eff.	YES	YES	YES	YES	YES
Boundary window (m)	∞	1,000	500	200	100
r2	0.940	0.942	0.944	0.951	0.961
N	1,082,777	656,353	338,982	125,209	56,640

Notes: Baseline model is as in column (3) of Table 2. Standard errors in parentheses and clustered on LE boundary x year effects. * p<0.1, ** p<0.05, *** p<0.01

Table E1: Varying boundary windows

	(1)	(2)	(3)
	log of sales price (in GBP)		
Imputed local broadband speed in	0.0251***	0.0252***	0.0253***
Mbit/s	(0.0016)	(0.0014)	(0.0014)
Speed^2	-0.0026***	-0.0026***	-0.0026***
	(0.0003)	(0.0002)	(0.0002)
4^{th} order distance poly.	YES	YES	YES
Control x year effects	YES	YES	YES
LE x year effects	YES	YES	YES
Excluded boundary window (m)	500	200	100
r2	0.936	0.933	0.933
N	743,795	957,568	1,026,137

Notes: Baseline model is as in column (6) of Table 2. Standard errors in parentheses and clustered on LE x year effects. * p<0.1, ** p<0.05, *** p<0.01

Table E2: Excluding boundary windows

Table E2 above presents estimates of eq. (5) excluding varying boundary window sizes. The models are otherwise identical to model (6) in Table 2. We stress that models (4) in Table E1 and (2) in Table E2 are mutually exclusive in terms of the considered samples. Yet, the results are highly consistent.

2. Functional form of the speed – LE distance relationship

To evaluate the sensitivity to the functional form of the speed distance relationship, we replicate our preferred models restricting the identifying variation to more simplistic functional forms of the speed – LE distance relationship in 2SLS estimation. We use two alternative predictors for speed and speed^2:

- Real speed spline (columns (1) and (3) in Table E3). This is an approximation of the speed-distance from LE relationship based on a linear spline function with a kink at 3 km (instead of a 4th order polynomial). This modification is meant to show that results do not really depend on the functional form of the engineering estimate.
- Technology x LE distance interactions (columns (2) and (4) in Table E3). We use dummies for technologies (ADSL, LLU, ADSL2+) and their interactions with distance from the LE. This model essentially corresponds to the baseline difference-in-differences model in Appendix D. It is supposed to show that the variation we need to generate our results really comes from technological upgrades and interactions with LE distance.

We implement the 2SLS method as follows: For each specification, the variables are first regressed against instruments and covariates (2 separate regressions, one for speed and another one for its square). The predicted values are then used instead of the actual values in the capitalization estimation of eq. (4). We stress here that we do not implement this imputation strategy to address endogeneity concerns. Our identification strategy takes care of this. We use this approach instead of the reduced form OLS (simply replacing speed and its square with the instruments) for convenience of interpretation. With this approach, we are able to alter the identifying variation while ensuring that the results are directly comparable to our baseline estimates.

Results are shown in Table E3. The real speed spline models (columns (1) and (3)) yield results that are extremely close to the baseline in the LE model. They are within a similar range although not as close in the BDD model. This is the expected result because by using an inferior functional form we introduce measurement error. It is not surprising that in the BDD, where we are generally further away from the LE and where we identify from a discontinuity that is supposed to be sharp, the attenuation bias is larger.

The technology x distance interaction models (columns (2) and (4))[37] produce very consistent estimates, which are close to the baseline results. They demonstrate that all the variation that is needed to generate our results comes from information about the timing of the upgrade and the distance of a property from the LE. The point estimates are slightly different, which is probably attributable to some measurement error being correlated with distance from the LE as we use a more simplistic spatial function. These are possibly the robustness checks that make most sense, since we impose just the very transparent structure for identification that we use in the difference-in-differences models, but have the advantage of estimating a marginal effect that directly speaks to the price-speed relationship.

	(1)	(2)	(3)	(4)
	log of sales price (in GBP)			
Imputed local broadband speed in Mbit/s	0.0207***	0.0285***	0.0247***	0.0288***
	(0.0052)	(0.0044)	(0.0016)	(0.0015)
Speed^2	-0.0009	-0.0035***	-0.0026***	-0.0036***
	(0.0013)	(0.0010)	(0.0003)	(0.0003)
4^{th} order distance poly.	YES	YES	YES	YES
LE effects	YES	YES	-	-
LE x year effects	-	-	YES	YES
Control x year effects	YES	YES	YES	YES
LE boundary x year effects	YES	YES	-	-
Boundary window (m)	200	200	-	-
Speed imputation	Real speed spline	Technology x distance interactions	Real speed spline	Technology x distance interactions
r2	0.9512	0.9512	0.9318	0.9318
N	125,209	125,209	1,082,777	1,082,777

Notes: Estimation method is 2SLS in all models. Predictors of the speed variables are described in the speed imputation row. Except for the speed variables, models in columns (1) and (2) are otherwise identical to Table 2, column (3), and models in columns (3), and (4) are otherwise identical to Table 2, column (6). Standard errors in parentheses are clustered on LE boundary x year effects in (1-2) and on LE x year cells in (3-4). * p<0.1, ** p<0.05, *** p<0.01.

Table E3: Restricted identifying variation

[37] Note that column (4) in Table E3 is the same as column (6) in Table 3 in the main text.

Appendix F: Heterogeneity

This section adds to section 4.3 in the main paper where we explore heterogeneity in the speed effect with respect to observable locational and individual characteristics.

1. Heterogeneity with respect to locational characteristics

In Table F1 we present the estimation results of three different versions of eq. (9). Column (1) allows for interactions between speed and income and urbanization (share of urban land within 1 square km). In column (2) we add interactions with various locational amenities. In column (3) we add further interactions with measures of local internet services.

	(1)		(2)		(3)	
	log of sales price (in GBP)					
Speed	0.026***	(0.001)	0.027***	(0.001)	0.027***	(0.001)
Speed^2	-0.003***	(0.000)	-0.003***	(0.000)	-0.003***	(0.000)
Speed x income	-0.001	(0.001)	-0.001	(0.001)	-0.001	(0.001)
Speed^2 x income	0.001***	(0.000)	0.000*	(0.000)	0.000*	(0.000)
Speed x urbanization	0.007***	(0.001)	0.004***	(0.001)	0.004***	(0.001)
Speed^2 x urbanization	-0.000	(0.000)	0.000	(0.000)	0.000	(0.000)
Speed x employment potential			0.002	(0.002)	0.004	(0.003)
Speed^2 x employment potential			-0.000	(0.000)	-0.001**	(0.000)
Speed x rail density			-0.003**	(0.001)	-0.003**	(0.001)
Speed^2 x rail density			0.000	(0.000)	0.000	(0.000)
Speed x school quality			0.003**	(0.001)	0.004***	(0.001)
Speed^2 x school quality			0.000	(0.000)	0.000	(0.000)
Speed x restaurant density			-0.001	(0.002)	-0.002	(0.002)
Speed^2 x restaurant density			-0.001	(0.000)	-0.000	(0.000)
Speed x retail density			0.008***	(0.002)	0.009***	(0.002)
Speed^2 x retail density			-0.001	(0.000)	-0.001	(0.000)
Speed x Amazon evening delivery					0.004**	(0.002)
Speed^2 x Amazon evening delivery					-0.000	(0.000)
Speed x # retailers dispatching					0.000	(0.001)
Speed^2 x # retailers dispatching					0.000	(0.000)
Speed x Uber					-0.005**	(0.002)
Speed^2 x Uber					0.001***	(0.000)
4^{th} order distance poly.	YES		YES		YES	
Control x year effects	YES		YES		YES	
LE x year effects	YES		YES		YES	
r2	0.932		0.932		0.932	
N	1,082,777		1,082,777		1,082,777	

Notes: Baseline model is column (6) in Table 2. All variables are normalized to have a zero mean and a standard deviation of one before being interacted with speed and speed^2. Standard errors in parentheses are clustered on LE x year effects. * p<0.1, ** p<0.05, *** p<0.01.

Table F1: Interaction effects

The strength of various interaction effects in our estimates suggests that significant heterogeneity exists in the speed effect. The income and urbanization interaction effects tend to become small as more interactions for amenities and local internet services are allowed for, which is in line with a concentration of these amenities and services in high income urban areas. Because the estimates of these interaction terms between the quadratic speed term and the relatively highly correlated locational variables are difficult to interpret, we first compute the marginal speed effect for every observation in our data and correlate it with the amenities considered in Figure F1 for ease of interpretability. We find that the marginal effect of real broadband speed tends to be larger in more urban areas with more amenities, suggesting that such areas are inhabited by households with a relatively high willingness-to-pay. Interestingly, the correlation is not similarly strong between the marginal speed effect and the median income of the local population.

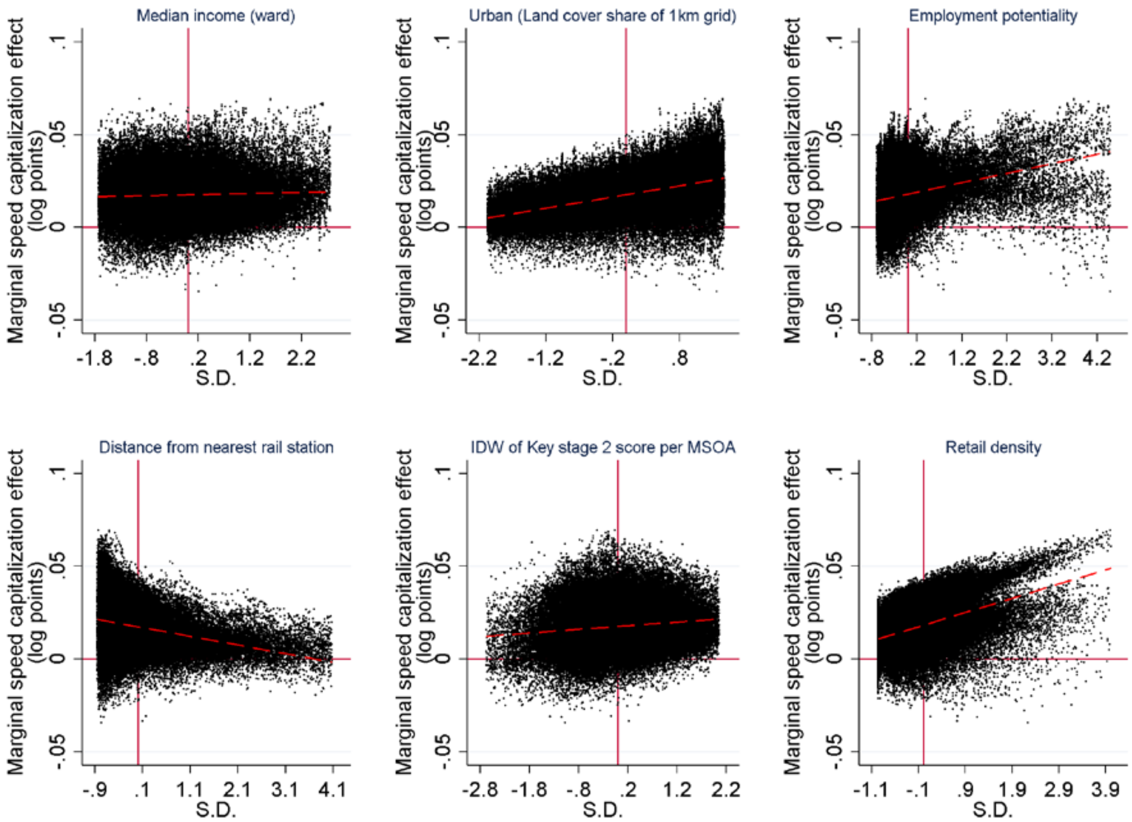

Notes: The figure shows unconditional correlations between the marginal speed effect derived from column (3) in Table F1 and selected locational characteristics where the distribution is normalized to have a mean of zero and a standard deviation of one. Observations below the 1st and above the 99th percentile in the distribution of the respective locational characteristic are omitted. To improve the quality of the presentation, a 1% random sample is plotted in each panel.

Figure F1: Interaction effects: Income, urbanization, and amenities

In Figure F2 we turn our attention to local internet services. We group properties by the number of major grocery retail chains that offer delivery of online orders. Then we compare the distribution of marginal speed effects across these groups. We find that the marginal speed effect increases in the number of grocery chains that dispatch online orders to a certain location, which is in line with the complementarity between broadband availability and local internet services suggested by Figure 7 in the main paper.

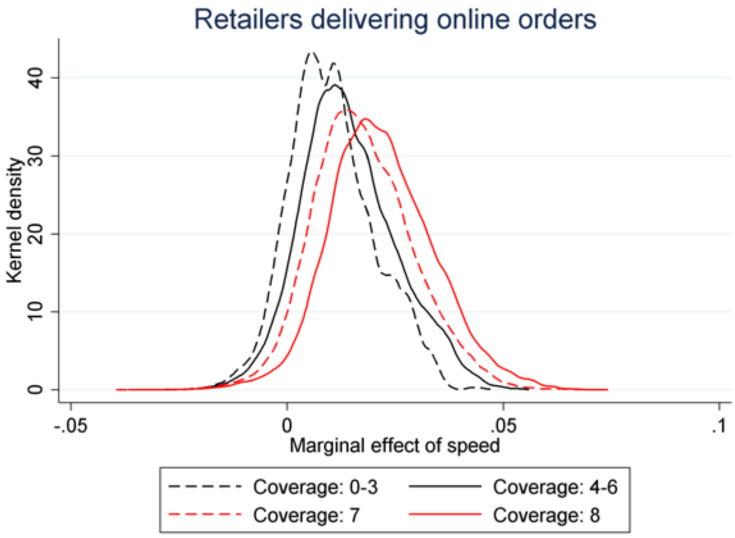

Notes: See Notes to Figure 7 in the main text.

Figure F2: Marginal speed effects: Retailers dispatching online orders

2. Heterogeneity with respect to individual characteristics

Table F2 below presents estimates of eq. (4) and (5) separately for samples of properties purchased by first-time buyers (columns 2 and 4) and other buyers (columns 1 and 3). The models are otherwise comparable to models (3) and (6) in Table 2. All estimates are close to the pooled baseline estimates (across all buyer types). Combined with the modest income effect discussed above these results suggest that economic and demographic individual characteristics may not be the primary determinants of the evident heterogeneity in the speed effects.

	(1)	(2)	(3)	(4)
	log of sales price (in GBP)			
Imputed local broadband	0.0298***	0.0246***	0.0251***	0.0258***
speed in Mbit/s	(0.0058)	(0.0075)	(0.0015)	(0.0020)
Speed^2	-0.0033**	-0.0020	-0.0025***	-0.0030***
	(0.0013)	(0.0017)	(0.0003)	(0.0003)
4^{th} order distance poly.	YES	YES	YES	YES
Control x year effects	YES	YES	YES	YES
LE effects	YES	YES	-	-
LE boundary x year effects	YES	YES	-	-
LE x year effects	-	-	YES	YES
Boundary window (m)	200	200	-	-
Buyer type	Non-FTB	FTB	Non-FTB	FTB
r2	0.9592	0.9584	0.9322	0.9290
N	76,196	49,013	720,392	362,385

Notes: Columns (1) and (2) are identical to column (3) in Table 2 except for separating the sample into first-time buyers (FTB) and non-FTB. Columns (3) and (4) are identical to column (6) in Table 2 except for separating the sample into FTB and non-FTB. Standard errors in parentheses are clustered on LE boundary x year effects in (1-2) and on LE x year cells in (3-4) * p<0.1, ** p<0.05, *** p<0.01.

Table F2: Capitalization effects by buyer type

Appendix G: Policy impact of the digital targets

Table G1 below reports the distribution of actual speeds by LE in our sample, organized by population decile. While this distribution is not exactly by density, as for the DCMS document, it is a good approximation, as faster broadband is typically deployed in more densely populated areas, while slower broadband exists in rural parts of the country. The distribution becomes our starting point for comparison with the speeds forecasted by the DCMS in 2020, presented in Table 5 in the main text. Notice that our speeds are *observed* actual speeds (see footnote 17), while the DCMS forecasts are in terms of the theoretical maximum speed attainable with a technology. Another reason for the large differences between our deciles and those in Table 5 is that our tests exclude cable subscribers, who generally connect to higher speeds. The "2010 target-equivalent" in our policy counterfactual is therefore an actual speed of 2.68 Mbit/s (or about 10 Mbit/s in terms of nominal speed).

Decile	1	**2**	3	4	5	6	7	8	9	10
Speed 2010 (in Mbit/s)	1.95	**2.68**	2.71	2.97	3.15	3.29	3.41	3.51	3.62	4.22
% of population with broadband connection	76.01	72.58	73.32	72.43	74.03	72.75	73.45	76.3	75.78	79.39

Table G1: Actual broadband speeds in England in 2010

We then follow the definitions of the EIB to attribute each LE to one of the three types of areas defined for the purpose of calculating costs (see Hätönen, 2011). According to these definitions, out of 22,925,211 English households, 85.89% are in LEs attributable to urban areas, 7.98% are in suburban areas, and the remaining 6.13% are in rural areas.

In Section 5, we report benefits at the household level to allow for comparisons of the speed upgrade and coverage upgrade. Hence, it does not actually matter how many people have to be connected when discussing values per household. Looking at the bigger picture, it is instead important to assess the aggregate benefits of the speed upgrade and coverage upgrade approach. For this, we proceed as follows. We add 10% to population covered, as, according to Ofcom, this is the percentage of people using mobile only for broadband purposes, which, therefore, will not need to be upgraded.[38] Then, we compute benefits in each LE, as reported in Table 6, and multiply those benefits per household by the number of households affected in that LE. We obtain aggregate total benefits of GBP 0.929bn for the speed upgrade and GBP 47.699bn for the coverage upgrade.

[38] See http://stakeholders.ofcom.org.uk/binaries/research/cmr/cmr13/UK_5.pdf.