Technological Change and Domestic Outsourcing*

Antonin Bergeaud Clément Malgouyres
Clément Mazet-Sonilhac Sara Signorelli

October 2023

Abstract

Does domestic outsourcing react to technological change? We study the staggered diffusion of broadband internet in France in the 2000s, and show that connected firms increased their outsourcing expenditures while decreasing the diversity of occupations they employ in-house. Meanwhile, employment in non-core occupations became increasingly concentrated in firms specializing in subcontracting services. Finally, we provide evidence that workers in high-skill occupations experienced salary gains from being outsourced, while workers in low-skill occupations lost out. Overall, we show that the deployment of new technologies stimulated domestic outsourcing in this context, with important implications for labor market inequality.

JEL classification: G14, G21, O33

Keywords: Broadband, Firm organization, Labor market, Outsourcing.

*Addresses: Bergeaud: HEC Paris, CEP-LSE and CEPR; Malgouyres: PSE, CREST and Institut des Politiques Publiques; Mazet-Sonilhac: Bocconi University, Finance Dept., Collège de France and IGIER; Signorelli: University of Amsterdam. We would like to thank Philippe Aghion, Adrien Bilal, Johannes Boehm, Mathilde Le Moigne and the seminar participants at College de France, Sciences Po, X-CREST, LAGV, JMA, AFSE, and RCEA for their helpful comments. Any errors are our own. Corresponding author: Antonin Bergeaud: 1 Rue de la Libération, 78350 Jouy-en-Josas, FR. Email: bergeaud@hec.fr

1 Introduction

The implications of international outsourcing (or offshoring), i.e., using a third-party firm based abroad to perform services that would otherwise be performed in-house by local employees, have been at the center of recurrent public debates and the focus of a large body of economic research and policy analysis.[1] More recently, empirical research has highlighted the pervasiveness of *domestic* outsourcing among developed economies (Dube and Kaplan, 2010; Goldschmidt and Schmieder, 2017; Bernhardt et al., 2016; Dorn et al., 2018) and has started to document its implications for the distribution of earnings. However, this nascent literature remains silent on the factors underlying these trends, particularly whether domestic outsourcing is primarily driven by changes in (labor market) institutions or by the adoption of new technologies.

In this paper, we look at the role played by innovation in Information and Communication Technology (ICT) in fostering the rise of domestic outsourcing. To do so, we estimate the impact associated with the diffusion of a new general-purpose technology, namely broadband internet (BI hereafter), on employment outsourcing by French firms. As explained in Abramovsky and Griffith (2006), general-purpose technologies, including ICT technologies, prompt important organizational changes within firms. In particular, by reducing communication and transaction costs, the diffusion of BI is likely to modify firms' optimal boundaries and affect the terms of their make-or-buy decisions (see, e.g., Lewis and Sappington, 1991; Garicano and Rossi-Hansberg, 2006; Bloom et al., 2014; Aghion et al., 2019b). Therefore, it is natural to assess BI's impact on the allocation of workers across businesses, particularly through an increase in domestic outsourcing.

We start by laying out a simple conceptual framework to discuss the different channels, identified by prior work, through which BI can impact domestic outsourcing. We

[1] See Hummels et al. (2018) for a review of recent empirical studies and Biscourp and Kramarz (2007) for a seminal work using French administrative data.

then use employee- and employer-level administrative data on the universe of French workers and firms from 1996 to 2007 to explore these links empirically. Specifically, we adopt an event study design, exploiting the staggered roll-out of the broadband infrastructure: the deployment of BI in France took nearly a decade to complete, from 1999 to 2007, and was particularly slow in the beginning. While cities connected earlier were typically larger and had higher population density, we exploit quasi-random variations in the timing of broadband connection across cities in the same province (*département* in French) with similar characteristics. Building on recent advances in the two-way fixed effects literature on staggered adoption contexts.[2], we adopt a stacked difference-in-differences approach (as developped in Cengiz et al., 2019).

Our results show that BI increased firms' expenditures in outsourcing services and the occupational segregation of establishments, as captured by a Herfindahl-Hirschman Index (HHI) of concentration. This effect is driven both by changes within existing establishments and changes in establishment composition, through entries and exits in the cities connected to the internet.[3] Interestingly, outsourcing affects not only low-skill workers such as cleaners or drivers, but also skilled professionals. To track these effects, we establish a list of sectors specializing in services to other firms, then we identify the occupations that are likely to be outsourced to such sectors.[4] We show that, following BI expansion, workers employed in "outsourceable occupations" in both the high- and low-skill segments become increasingly concentrated in establishments specializing in these particular services, and thus less likely to be employed in-house by establishments active in other sectors. Finally, we provide suggestive

[2] See e.g. Callaway and Sant'Anna (2020); De Chaisemartin and d'Haultfoeuille (2020); Goodman-Bacon (2018); Borusyak et al. (2021).

[3] This evidence is in line with the trends described by Godechot et al. (2020), who, based on administrative data in several high-income countries, document a steep increase in both earnings and occupational segregation at work. Especially with regard to occupation, this result is particularly pronounced in France.

[4] A detailed description of these definitions is reported in Appendix B. The definition of low-skill outsourcing follows Goldschmidt and Schmieder (2017) closely and includes cleaning, security, driving, and logistics. For high-skill outsourcing, we select the two largest industry categories that provide professional services to other firms: IT and consulting (which includes strategy consulting, HR and advertising).

evidence that the impact of domestic outsourcing on wages is heterogeneous across skills: high-skill workers experience wage increases after outsourcing while low-skill workers see their wages decrease, implying that outsourcing tends to widen preexisting earning disparities.

The analysis in this paper contributes to several strands of literature. First, we relate to recent empirical analyses exploiting BI diffusion as a plausibly exogenous technological and informational shock. In France, Malgouyres et al. (2021) show that the BI expansion was associated with an increase in imports among affected firms.[5]

More closely related to our study, Akerman et al. (2015) evaluate the skill bias generated by BI in Norway. They find that broadband availability increases both firm productivity and the high-skill wage premium. We contribute to this literature by diving deeper into the concrete mechanisms through which skill bias occurs. In particular, we show that it arises partly through the reorganization of firms and the reallocation of high and low-skill workers across firms, not just through a shift in the firm-level production function. Given frictional labor markets, the allocation of workers to firms has important implications for the distribution of earnings (Song et al., 2019). Consistent with this point, we show that these reorganizations affect heterogeneous workers in a way that amplifies preexisting wage inequalities.

Second, our paper is linked to the literature on domestic outsourcing.[6] Goldschmidt and Schmieder (2017) show how Germany has experienced an explosion in the domestic outsourcing of low-skill occupations since the early 1990s, which resulted in wage reductions for outsourced workers, mainly driven by the loss of firm-specific

[5] Consistently, Akerman et al. (2018) show that BI narrowed the role of distance in explaining bilateral trade in Norway. In the United Kingdom, DeStefano et al. (2018) find that a similar shock affected firm size but not firm productivity, which suggests that local institutions matter in evaluating the effect of broadband expansion (see also Haller and Lyons, 2015). On the contrary, Bertschek et al. (2013) finds a positive causal effect of BI on innovation and productivity in Germany (see also Fabling and Grimes, 2021; Grimes et al., 2012).

[6] Weil (2014) describes how the nature of work has changed in the 21st century as more large companies have switched to a "fissured workplace" business model.

rents.[7] Bilal and Lhuillier (2021) consider the aggregate implications of domestic outsourcing in terms of wages and productivity. Other work has further highlighted a possible link between technology and outsourcing.[8]

We contribute to this literature in three main aspects. First, we exploit a natural experiment that allows us to test the causality of the relation between a specific technology—BI —and domestic outsourcing. Second, we extend the spectrum of affected workers to high-skill occupations, which enables us to explore the heterogeneity of the wage effect across skill groups.[9] Third, we capture margins of adjustment largely ignored by previous studies. Although most earlier work focuses on individual transitions from in-house to outsourced activities (Dube and Kaplan, 2010), or on plant-level outsourcing events (Goldschmidt and Schmieder, 2017), the rise in domestic outsourcing is likely also driven by newly-created firms using outsourcing services more intensely than firms exiting the market, and by the concentration of new entrants in the service-providing sectors. Our city-level approach allows us to capture such margins, and comparing them with the results obtained from a plant-level analysis on incumbents clarifies the magnitude of the composition effects. In most cases, we find that city-level estimates are larger than plant-level ones, consistent with the notion that firm entry is an amplifier of the effect of BI on domestic outsourcing.

The rest of the paper is organized as follows. Section 2 briefly discusses the mechanisms through which a technological change such as BI can be related to an increase

[7]They focus on four emblematic activities— cleaning, logistics, security, and food service— and introduce an innovative measure to capture the outsourcing of such tasks. Drenik et al. (2020) quantify the resulting loss in wages and show that outsourced workers share about half the share of the firm-specific wage premium that insiders earn.

[8]Using data on US firms, Fort (2017) shows that technology lowers the cost of coordination within a firm and ultimately increases the fragmentation of production. Cortes and Salvatori (2019) show that firms in the UK have become increasingly specialized over the past 20 years, partly driven by an increase in domestic outsourcing for high-skill, non-routine cognitive tasks. Aghion et al. (2019a) characterize so-called "good jobs" that are protected from outsourcing as technology advances, even among less-educated workers.

[9]Abramovsky et al. (2017) discuss the evolution of high-skill occupations' *offshoring*. The literature has also identified occupations that are more likely to be domestically outsourced to the growing business service sector, including some high-skill occupations such as advertisers, accountants, IT specialists, and legal professionals (Ono, 2003; Berlingieri, 2014; Goldschmidt and Schmieder, 2017).

in domestic outsourcing. Section 3 details the data and our empirical methodology, while Section 4 presents results on the effect of BI on outsourcing. Finally, Section 5 looks more closely at the consequence of outsourcing on workers' wage trajectories, and Section 6 concludes the paper.

2 Technology-led domestic outsourcing?

In this section, we discuss the mechanisms through which access to broadband internet can affect a firm's decision to outsource. Abramovsky and Griffith (2006) argue that ICT adoption at the firm level is likely to lower the adjustment and monitoring costs of outsourcing. BI could act in the same way, by facilitating the exchange of information between regular firms and their various subcontractors.[10] Given that outsourcing costs are heterogeneous across firms, for instance because of differences in product complexity (Tadelis, 2002), we expect BI to lead to a greater proportion of firms finding it financially beneficial to outsource certain activities rather than produce them in-house.[11]

Abraham and Taylor (1996) distinguish three additional motives underpinning a firm's decision to outsource its labor: (i) the realization of wage and benefit savings, (ii) access to the specialized skills possessed by subcontractors, and (iii) the management of volatility in output demand. We focus our discussion on the first two motives. The first is connected to a recent literature that emphasizes how labor market frictions lead to some rent-sharing between firms and in-house workers (Card et al., 2018; Lamadon

[10] In fact, as we will show in Figure II, BI caused an increase in firm-level IT hardware intensity. Accordingly, the BI shock might be operating through the very mechanism conceptually and descriptively highlighted by Abramovsky and Griffith (2006). See subsection "Broadband availability versus broadband adoption" in Section 3.3 for more details on this result.

[11] It is important to highlight the distinction between temporary agency workers and outsourced workers. Temporary agency workers are employed by a staffing agency and assigned to a client firm, where they perform tasks as in-house employees (e.g., covering the front desk for an individual on sick leave). The utilization of temporary agency workers is subject to regulatory constraints and limitations. In contrast, outsourced workers, which are the focus of this paper, are provided by a specialized firm and subject to minimal control by the client company.

et al., 2019; Kline et al., 2019). This creates an incentive to outsource some occupations to outside contractors, as the market price for these services is independent of the firm's rents.[12] This incentive grows as firms become more productive or enjoy larger rents. In line with Akerman et al. (2015), a first effect of BI is to increase firm-level productivity, triggering more outsourcing as firms expand in imperfect labor markets. The relevance of this motive is consistent with a key empirical regularity: the positive relationship between firm size (measured by revenue) and the propensity to outsource, which is shown in Figure I.

FIGURE I. Correlation between outsourcing intensity and firm sales

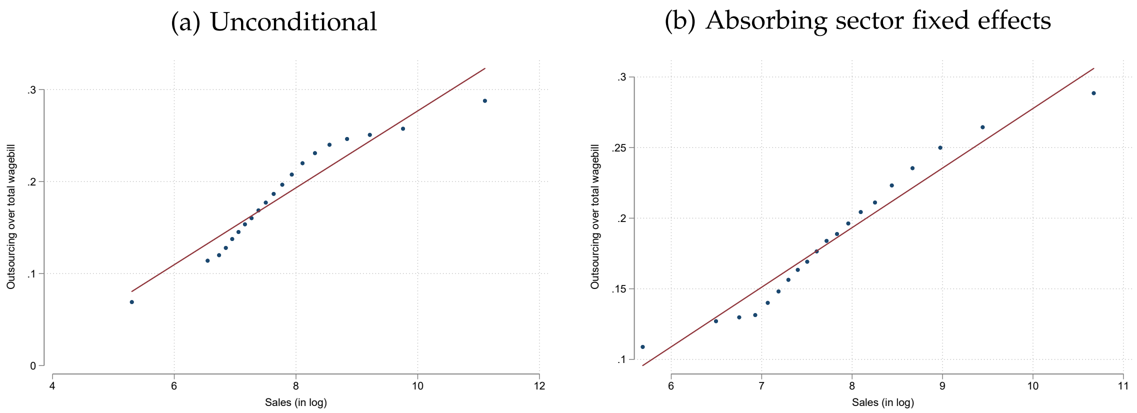

Notes: The figure presents the binned scatter plot for the correlation between firm outsourcing intensity and the log of firm sales. The left panel reports the row correlation and the right panel reports the correlation after the 5-digit sector fixed effect is absorbed.

The second motive for outsourcing is linked to subcontractors' ability to provide cheaper quality-adjusted services, which derives from their access to more specialized skills thanks to economies of scale (Bartel et al., 2005). Subcontracting firms offer their services to multiple clients. Accordingly, their scale of production is likely larger than those of their individual clients if those clients were to produce the service in-house. To the extent that BI boosts firms' productivity, it will increase the size of the market for outsourcing services, which in turn allows for greater economies of scales, thus magnifying the incentives to outsource.

[12]Even if there is some rent-sharing with respect to outsourced workers, to the extent that this rent-sharing is weaker than with respect to in-house workers— as shown by Drenik et al. (2020) in the case of temp workers— firms will remain incentivized to outsource part of their production.

Finally, in the case of IT services, there is one additional motive to be noted. BI is likely to boost directly the efficiency of IT subcontractors, driving their marginal revenue productivity further up. For the IT-intensive outsourcing sector, the BI shock corresponds to both an increase in efficiency and a reduction in cost. If barriers to entry are low, new IT service firms will enter the market and put a downward pressure on prices, increasing regular firms' incentives to outsource.[13]

In Appendix E, we sketch a simple partial equilibrium model to highlight a non-trivial mechanism related to the first motive described by Abraham and Taylor (1996): wage and benefit savings in presence of labor market frictions. We show how a positive, Hicks-neutral productivity shock can lead firms to outsource a (non-random) part of their workforce. This effect relies on the assumption that firms possess some wage-setting power when hiring in-house workers, but they are price-takers when they buy outsourcing services.[14] This difference implies that the relative cost of an in-house workforce versus outsourcing increases as the firm's optimal size goes up, following for instance a surge in productivity. This mechanism generates the positive correlation between outsourcing intensity and size displayed in Figure I. In this context, a technological shock such as broadband expansion, which boosts firms' productivity, leads to more outsourcing. Importantly, the model sketched in Appendix E generates two additional testable predictions that we bring to the data. First, BI increases outsourcing intensity, measured as the ratio of outsourcing costs to in-house workers' payroll. Second, under a broad set of conditions, it also leads firms to become more segregated in the type of workers they employ, by hiring a smaller set of occupations

[13] Accordingly, the producer price index (PPI) of the broader sector that includes most outsourcing service firms (unfortunately, PPI values are only available at the 2-digit sector level) over the PPI of the total economy declined by about 0.05 between 2000 and 2007.

[14] This assumption is in line with empirical evidence obtained in the case of temporary workers— a specific instance of domestic outsourcing (Drenik et al., 2020). Firms face imperfectly elastic labor supply curves and will accordingly need to raise wages in order to expand their in-house workforce. Imperfect elasticity in labor supply is generated by idiosyncratic workers' tastes for the amenities offered by firms (such as working conditions, commute, and corporate culture), as is standard in the recent literature on monopsony (Card et al., 2018; Lamadon et al., 2019). Given asymmetric information on workers' heterogeneous preferences for jobs, firms are unable to perfectly discriminate and fully price these amenities into individual-specific wages.

in-house.

To sum up, a succinct review of potential mechanisms reveals that there are several theoretical reasons why BI access could boost outsourcing, including lower transaction costs, how firms expand in presence of imperfect labor markets, and rising productivity among IT subcontractors. We now turn to the empirical analysis to assess the magnitude of this effect.

3 Data and Empirical Strategy

3.1 Data

Administrative data on labor market outcomes. Our data comes from two main administrative sources. The first is the matched employer-employee dataset covering all workers based in France since 1994, the (*Déclaration annuelle des données sociales* or DADS). This dataset constitutes the primary source used to compute payroll taxes and gives detailed individual information including salary, hours worked, occupation, age, gender, and an identifier showing the employing firm and establishment. It is not possible to follow these workers over time, except for a random sample of 1 out of every 24 workers, which we refer to as the DADS panel. In this data, we also lack information on the level of education attained. Therefore, we always define the skilled nature of jobs based on the occupation codes, which are highly correlated with educational attainment in the French context.

Our second main data source is the administrative financial records curated by the Bank of France and called *Fichier inter-Bancaire des Entreprises* (or FiBEn). FiBEn includes firms with a yearly turnover of 750,000 euros or higher, and has the advantage of recording outsourcing expenditures as a separate item.[15] Contrary to the DADS,

[15] In contrast, in the more widely used *FIchier Complet Unifié de Suse* (FICUS), the value of outsourcing expenditures is included within a broader category of spending, making it impossible to isolate the impact of broadband internet on this particular sub-category.

FiBEn does not provide information at the establishment level, but only at the more aggregate firm level.[16] We clean these data and restrict our sample to the French private sector. Appendix B provides details on how we depart from the raw dataset.

We construct several measures as proxies for domestic outsourcing. The most direct one is firm-level outsourcing expenditures, scaled by the in-house wage bill.[17] To gain a broader sense of how BI affects the segregation of occupations across establishments, we construct a measure of establishment occupational concentration, based on an HHI computed using shares in the total wage bill.[18] Finally, we define typical outsourcing sectors and outsourceable occupations as follows, consistently with Goldschmidt and Schmieder (2017):

- **Outsourceable occupations**[19]: Non-core occupations that are likely to be outsourced, off-shored, or subcontracted. We divide them into high-skill outsourceable occupations (IT specialists and consultants) and low-skill outsourceable occupations (cleaners, security workers, drivers, and logistics workers). The full list of occupations and their definitions is detailed in Appendix B.

- **Outsourcing sectors**: Administrative and business support service sectors. High-

[16]Our preferred model for the outsourcing expenditure results keeps multi-establishment firms and assigns outsourcing intensity to the location of the headquarters, since information on the other plant locations is not available.

[17]The available information on outsourcing expenditure does not allow us to distinguish between international outsourcing (offshoring) and domestic outsourcing. We expect domestic outsourcing to represent the first order response as international trade in services remains fairly rare and a lot of the services that we consider (especially in the low-skill outsourcing sector) are inherently non-tradable activities (cleaning and security for instance). Moreover, using survey data from 2016 (INSEE Survey on Global value chains of SME), 7.4% of firms report doing domestic outsourcing versus only 2.4% for offshoring, an odds ratio of about 3.

[18]The concentration of occupations within each establishment i is computed with the following formula: $HHI_i = \sum_{o=1}^{18} \psi_{oi}^2$, where o indexes each of the 18 occupational categories that compose the French private sector at the 2-digit level (also known as the PCS classification), and ψ_{oi} is the share of the wage bill of the establishment attributed to workers employed in each occupation. At the city level c, we compute the average establishment concentration as follows: $HHI_c = \sum_{i=1}^{N} \omega_{ic} HHI_{ic}$, where ω_{ic} is the share of employment in the city accounted for by each establishment i.

[19]Throughout the paper, we make the assumption that one occupation is a fixed bundle of tasks. This is admittedly a simplification, since the literature has shown that the task content of jobs evolves over time, including evolution due to technological change (e.g. Spitz-Oener, 2006, Deming and Noray, 2020). We are obliged to make this assumption due to the limits of the administrative data, which does not allow us to observe tasks beyond their occupational dimension.

skill outsourcing sectors: IT and consulting sectors. Low-skill outsourcing sectors: cleaning, security, driving, and logistics services. The full list of sectors and their definitions is detailed in Appendix B.

At the city level, we look at the expansion of the outsourcing sectors by measuring their share of overall employment, and we look at shifts towards these sectors by measuring the share of workers in outsourceable occupations that are employed by them and not elsewhere. At the individual level, we look at the mobility of workers across firms and focus on moves towards the outsourcing sectors.

Data on broadband expansion ADSL (Asymmetric Digital Subscriber Line) is a data communication technology that enables fast data transmission over copper telephone lines. The eligibility for ADSL depends on the distance between the final customer and a Local Exchange unit (LE), since the quality of the signal decreases as it is routed over the copper lines. Our main information of interest documents the date when each LE was upgraded in mainland France.[20] While we do not have data on firm-level use of BI, suggestive evidence suggests that ADSL was widely adopted by firms once it became available to them.[21] Additionally, we obtained data from the regulatory agency (ARCEP) on the geographical coverage of each local exchange unit. Each city in France is partitioned into census blocks, and the data document the area of each census block covered by a given local exchange unit. Combining both datasets, we construct a continuous measure of broadband access for city i at year t. This measure, which we denote as $\widetilde{Z}_{it}$, is a time-weighted percentage of the area covered by BI in

[20] See Appendix C for a more detailed account of the ADSL expansion. The data was collected through the official website by Malgouyres et al. (2021) and validated manually.

[21] A 2016 survey by the French regulator shows that in that year, 73% of small and medium enterprises used ADSL technology (Arcep, 2016). The large take-up reflects the fact that ADSL was a massive improvement in terms of speed, as well as in terms of connection cost and time. Additionally, we find that firms increased their investment in IT hardware in response to the the BI shock. See subsection "Broadband availability versus broadband adoption" in Section 3.3 for more details on this result.

city i. It is formally defined as:

$$\widetilde{Z}_{it} = \sum_{b \in i} D_{b,t} \frac{\mathcal{A}_{b,t}}{\sum_{b' \in i} \mathcal{A}_{b',t}}, \tag{1}$$

where $b \in i$ denotes the census blocks included in city i and $D_{b,t}$ is the share of the days of year t with BI access in b. Finally, $\mathcal{A}_{b,t}$ denotes the area covered by census block b. The variable $\widetilde{Z}_{it}$ is continuous with support between 0 and 1.[22] In most of our study, we discretize $\tilde{Z}$ by defining the year of treatment as the first year where $\tilde{Z}$ becomes positive. We denote as $C_{i,t}$ the corresponding discrete (binary) variable, and we refer to it as the "year of connection." This transformation results in little loss of information, given the underlying distribution of the continuous variable, and allows us to implement a transparent before/after comparison through an event study design. Figure A3 in Appendix A shows the evolution of the average $\tilde{Z}$ within cities after the year of first connection. On average, the first year is characterized by a degree of connection of 30%, which jumps to 60% the following year and quickly reaches 80%.

Table I reports the summary statistics for the main outcomes of interest. Panel A describes the city-level data while Panel B describes the establishment level data. On average, firms spend on outsourcing 30% of the amount they spend on in-house workers. The high- and low-skill outsourcing sectors account for 2% and 11% of employment, respectively, and 8% (25%) of employment in high-skill (low-skill) outsourceable occupations is located within their respective outsourcing sector.

[22] $\widetilde{Z}_{it}$ will be equal to one if all of its areas had BI access for the entire year. It will be equal to 1/2 if the entire city had access to broadband over half the year t. In practice, however, it is very strongly concentrated at 0 and 1, with very few intermediate observations; see Figure A2 in Appendix A.

TABLE I. Summary Statistics

	mean	(sd)
Panel A: City level		
Share high skill workers	0.07	(0.07)
Occupational concentration (HHI)	0.38	(0.16)
Value added over wage bill (2010 euros)	2.08	(1.84)
Outsourcing exp. over wage bill (2010 euros)	0.31	(0.58)
Share of empl. in high-skill outs. services	0.02	(0.09)
Sh. outsourceable workers in HS outs. Services	0.08	(0.22)
Share of empl. in low-skill outs. services	0.11	(0.24)
Sh. outsourceable workers in LS outs. Services	0.25	(0.38)
Average establishment size	77	(263)
N. of establishments in city	13.2	(59.7)
Panel B: Establishment level		
Share high skill workers	0.11	(0.17)
Occupational concentration (HHI)	0.38	(0.19)
Value added over wage bill (2010 euros)	2.19	(5.78)
Outsourcing exp. over wage bill (2010 euros)	0.40	(1.13)

Notes : This table presents the summary statistics of the main variables used in the regression. Averages and standard deviations are computed over the full period from 1997 to 2007. Observations: 172,132 for Panel A and 2,439,340 for Panel B

3.2 The diffusion of Broadband Internet in France

As evidenced by Malgouyres et al. (2021), the deployment of BI beyond France's largest cities began very slowly at the dawn of the 2000s, due to multiple reasons linked to the country's historical telecom supplier.[23] It then accelerated between 2004 and 2007, thanks to the involvement of regional governments who favored competition among providers. Table II shows the percentage of cities, establishments, and workers that were connected each year. In 2000, only 2% of cities were connected, although this corresponds to a much larger share of workers and establishments (respectively, 25% and 22%). By 2003, 80% of workers and 76% of establishments were connected. At the end of 2005, 96% of French workers and 80% of cities were covered.

[23]First, France Télécom, the monopolistic telecom supplier, was uncertain regarding the future regulations on the wholesale price that it was going to be able to charge. Second, at the same time that France Télécom had to invest massively in upgrading its LEs to ADSL, it went through a debt crisis that ended in 2002 with what was essentially a government bailout. Urged on by the government—which increased its stake in the firm during the bailout —in 2003, France Télécom pledged to cover 90% of the French (mainland) population by the end of 2005.

TABLE II. Year of connection by city

Year of Connection	Number connected (in % of total)			
	Cities	Workers	Establishments	Population
2000	2.1	25.0	22.2	18.5
2001	6.6	35.7	34.7	28.9
2002	8.4	19.3	19.6	18.4
2003	12.4	6.7	7.8	9.5
2004	18.4	5.0	5.8	8.4
2005	23.0	4.4	5.4	8.5
2006	18.6	2.2	2.8	5.3
2007	8.8	1.6	1.7	2.5

Notes: All values are taken in 1999. The sum of percentages in a column is different from 100 because a small number of cities are not connected to the ADSL in 2007. The number of establishments and workers is based on our final sample (therefore following our cleaning and selection procedures).

Given that our main effects of interest are identified out of the gradual diffusion of BI across cities, addressing the endogeneity of the decision to "treat" one LE before another deserves special consideration. Malgouyres et al. (2021) describe how broadband expansion occurred to maximize population coverage with no special consideration for economic potential (see Appendix C for more detail). Our identification strategy relies on the assumption that, within a given department,[24] the coverage of cities was mostly determined by city population density— which is almost fixed over time and can be controlled for— and did not take into account underlying local trends in productivity or propensity to outsource activities. As a result, conditional on time and city fixed effects, we consider the variation in broadband access to be as good as random.[25] In the next section, we detail how we test the plausibility of this assumption.

[24]In France, a department *(département)* is a type of administrative division. There are 94 department located in mainland France, with an average population of 700,000 in 2019. On average, each department encompasses around 300 cities, which will serve as the primary unit of analysis in this paper. The map at the top of Figure A1 provides a visual representation of the size of both departments and cities in France.

[25]See Malgouyres et al. (2021) for an in-depth discussion of the determinants of the geography of broadband expansion.

3.3 Empirical Strategy

To identify the causal effect of BI, we exploit the fact that the dissemination of the ADSL technology across cities was staggered over a period of nearly 10 years, from 1999 to 2007. While the timing of diffusion is likely to be correlated with the size and density of the city, after controlling for these elements we can exploit quasi-random variation in access, across cities of similar densities that are located within the same department. Formally, we use the panel of cities c observed for each year t from 1997 to 2007 to run the following econometric model:

$$Y_{c,t} = \sum_{\tau=-k}^{k'} \alpha_\tau \mathbb{1}\{t = t_c + \tau\} + \boldsymbol{\gamma_t X_{c,t}} + \nu_c + \varepsilon_{c,t} \tag{2}$$

where Y is the outcome of interest, t_c is the year of BI arrival in a city c, and X is a vector of control variables. X contains a time-unvarying measure of the density of the city in 1999, interacted with year dummies and a set of department-year fixed effects. Finally, we add city-level fixed effects such that the coefficients α_τ can be interpreted as changes within a given city resulting from the arrival of BI, compared to cities within the same department and with similar densities that have yet to be connected to BI. The regression is run over the sample of cities that have more than 100 inhabitants at the beginning of the period, to avoid capturing an effect driven by small villages, and over the years 1997-2007, which include the full period of BI expansion. To avoid capturing endogenous changes in establishment location as a response to BI diffusion, we fix the city as the one where each establishment appears for the first time in our data, and we keep it constant throughout the period.

Recent literature underlines the caveats of using two-way fixed effects models in staggered adoption contexts (Callaway and Sant'Anna, 2020; De Chaisemartin and d'Haultfoeuille, 2020; Goodman-Bacon, 2018; Borusyak et al., 2021). The main issue resides in the fact that α_τ for the post-treatment period is partly estimated using previously treated observations as controls. This might introduce biases in the presence of dynamic or heterogeneous treatment effects across cohorts. Given that in our context there is no pure control, since all cities were treated by the end of the period, we adopt the stacked difference-in-differences approach employed by Vannutelli (2020), which is also similar to Cengiz et al. (2019) and Deshpande and Li (2019).[26] The latter consists of constructing a *rolling control group* for each treated cohort. In practice, for each treated cohort between 2000 and 2006, we construct a separate sample where we define the time to treatment relative to that specific cohort, and where all the cohorts treated in later years serve as controls. The observations within the control cohorts are only considered for the period preceding their respective treatment. The cohort receiving BI in 2007 serves as a pure control, since there is no cohort treated afterwards to create an additional sample. Finally, we append the 8 samples constructed using this method to run the following regression model:

$$Y_{c,s,t} = \alpha_0 Treat_{c,s} + \sum_{\tau=-k}^{k'} \alpha_\tau D^\tau \times Treat_{c,s} + \sum_{\tau=-k}^{k'} \beta_\tau D^\tau + \boldsymbol{\gamma_t X_{c,t}} + \nu_c + \pi_s + \varepsilon_{c,s,t} \quad (3)$$

Where $Treat_{c,s}$ takes the value of 1 if city c is treated in sample s. This parameter is identified despite the city fixed effects ν_c, because the same city appears in multiple samples with both treated and control status. D^τ are dummies for time relative to treatment (equal to $\mathbb{1}\{t = t_c + \tau\}$ in equation (2)) and α_τ identify the pre- and post-treatment dynamic effects. We also add fixed effects for each of the samples stacked (π_s). Standard errors are clustered at the department level, which also accounts for the error correlation generated by the repeated appearance of the same cities across different samples. This constitutes our preferred strategy that we apply throughout the paper, but in Appendix A we also report the results obtained from the standard

[26] Baker et al. (2021) show that the stacked difference-in-differences method gives similar results to the method proposed by Callaway and Sant'Anna (2020) in a variety of corporate finance applications. Gardner (2022) shows that this approach estimates a weighted average of the cohort-period specific average treatment effect on the treated where the weights are a positive function of the number of treated units and the variance of treatment within each stacked event.

two-way fixed effects model presented in equation (2), for comparison.[27] Finally, in the main graphs of the paper we also show the robustness of our results to controlling for the propensity of cities to adopt BI early based on observable baseline characteristics, which we describe in detail below.

Tests of the identification assumptions First, our approach assumes that the timing of BI arrival within a department is as good as random, conditional on trends explained by city density, and on time and city fixed effects. The latter might be violated if BI is deployed first in cities that were already on a positive productivity trend or that are larger in terms of economic activity (conditional on city density), in cities with a specific industry composition, or in cities going through changes in the political party in power. To test the validity of this assumption, we run alternative specifications adding controls for each one of these elements, constructed as follows:

1. **Productivity trends**: growth in labor productivity between 1996 and 1998, measured as average value added per worker in the city, interacted with year fixed effects.[28]

2. **Size of economic activity**: number of establishments in the city in 1997, interacted with year fixed effects.

3. **Industry composition**: shares of employment in 11 broad economic sectors in 1997, interacted with year fixed effects.

4. **Changes in politics**: share of votes for left-wing candidates at the presidential

[27] In the standard TWFE model it is necessary to exclude two dummies from the regression, respectively for $\tau = -1$ and $\tau = -6$, to avoid multi-collinearity and to identify the fully dynamic underlying data generating process in this staggered design with year and individual fixed effects (Borusyak et al., 2021). However, the stacked approach in our main analysis does not require this restriction.

[28] The value-added information is sourced from the FICUS administrative dataset, which encompasses all firms (except financial services firms) and does not enforce a minimum yearly turnover of 750,000 euros, unlike FiBEn. Although this limitation affects the number of firms (approximately 30% of the FiBEN sample), FiBEn's coverage is generally considered extensive, representing about 70% of the total workforce and value-added data.

election of 1995 and a dummy for change in the political majority between 1995 and 2002, both interacted with a full set of year fixed effects.[29]

Results are presented in Appendix A.3. Given that adding multiple controls quickly runs into the course of dimensionality, we also compute a propensity score for early treatment based on city baseline characteristics, which we interact with year fixed effects and which we include as controls.[30] Results are presented alongside the main coefficients in Section 4 to highlight their robustness. In addition, we use the propensity score to perform placebo tests, where late receivers of broadband (2004 to 2007) are split into high and low propensity to be early receivers (HPLR and LPLR). We then attribute a fictitious treatment date to each city among the late receivers by simply taking the actual date minus 4 years. Finally, we run a regular event-study regression comparing HPLR to LPLR on a sample period, where none of them is actually treated (1997 to 2003) and where therefore we expect no effect of BI. These placebo tests are presented in Appendix A.3.

Our approach further assumes that cities located in the same department and receiving BI later in time are not affected by neighboring areas adopting the technology earlier. This hypothesis would be violated if the arrival of BI were associated with important changes in workers' mobility across cities. To test the plausibility of this assumption, we show that our results are robust to: i) running the analysis at the broader commuting zone (CZ) level instead of the city level,[31] ii) including commut-

[29]These variables are proxies for the political orientation of local politicians, since candidates in local elections often do not have an explicit political affiliation. We categorize all the national candidates that ran in the 1995 and 2002 presidential elections as either left-wing or right-wing. Left-wing candidates include Lionel Jospin, Robert Hue, Arlette Laguiller, Dominique Voynet, Jeacques Cheminade, Jean-Pierre Chevènement, Noël Mamère, Olivier Besancenot, Christiane Taubira, and Daniel Gluckstein. Right-wing candidates include Jacques Chirac, Edouard Balladur, Jean-Marie Le Pen, Philippe De Villiers, François Bayrou, Jean Saint-Josse, Alain Madelin, Bruno Mégret, Christine Boutin, and Corinne Lepage.

[30]We define early adopters as the cities receiving broadband in the first 3 cohorts (1999 to 2001). We then estimate a linear probability model where the predictors are the number of establishments in the city, the sectoral composition, the average level of productivity of all firms (measured before 1999), and the productivity growth observed between 1996 and 1998.

[31]Commuting Zones (*"Zones d'emploi"*) or CZs are smaller entities than departments but larger than

ing zone × year fixed effects instead of department × year fixed effects,[32] and iii) adding controls for the BI coverage of other cities within the department to the main specification. All these robustness checks are presented in Appendix A.3. In addition, our individual-level results show that the additional worker mobility generated by BI across cities is small enough to be considered negligible from the point of view of the labor markets of neighboring cities (results shown in Table III and discussed in section 4.2).

Finally, given that our main analysis applies a binary treatment indicator based on the year BI first arrives in the city, one might wonder whether the dynamic effects recovered are driven by the effect of the treatment becoming stronger over time or by the BI coverage of the city becoming larger over time (Figure A2 in Appendix A shows that coverage increases in the first two years following the first BI connection). We show that the dynamic effects are not confounded by changes in coverage, by providing results using the continuous variable $\tilde{Z}_{it}$ instead of the binary treatment. All of these robustness checks are presented in Appendix A.3.

Broadband availability versus broadband adoption. Our paper focuses on a reduced-form estimation in the sense that it captures the effect of local broadband availability on firm behavior and not the effect of firm-level adoption of the technology. Accordingly, our estimates reflect the combination of two effects: the effect of broadband availability on its adoption by firms and the effect of firm-level adoption on firm behavior. Given that we do not have data on firm-level adoption, we cannot disentangle them. However, we view the reduced-form coefficients as policy-relevant, since it is straightforward to increase local broadband availability through public policy, while

cities. They correspond to labor market areas. There are 297 CZs in France, which on average contain 117 cities. See Figure 1(b) for an illustration.

[32]In the presence of strong geographic spill-overs, restricting the control group within a commuting zone would magnify the bias and thus is expected to generate different coefficients.

manipulating adoption involves more intricate interventions.[33] Moreover, building on newly available data for a subset of the sample period (1998-2007), we are able to construct an indicator of IT hardware intensity defined as the (book) value of IT capital over total capital at the firm-level.[34] We average this variable at the city level and we assess, using the baseline specification described in Equation (3), whether it is positively affected by BI availability. As shown in Figure II, we find that this measure increases significantly following BI expansion. The increase is about 3 percentage points four years after connection (from an average pre-broadband level of 7%). Naturally, this is not a direct measure of internet adoption but rather a measure of investment in equipment likely to be complementary to internet usage. While we do not view this exercise as a direct first-stage analysis, we consider this result as supporting the notion that our reduced-form empirical approach effectively captures changes in the behavior of firms with respect to internet use and IT technology more broadly. Moreover, to the extent that ICT investment is associated with a higher propensity to outsource (Abramovsky and Griffith, 2006), this result implies that broadband expansion is likely to boost outsourcing through this channel.

[33]See e.g. Andrews et al. (2018) for a cross-country empirical investigation of the numerous determinants of ICT adoption by firms and their complex interaction.

[34]Recently, Lashkari et al. (2022) use this dataset to study the link between IT hardware investment and returns to scale.

FIGURE II. IT hardware capital and BI

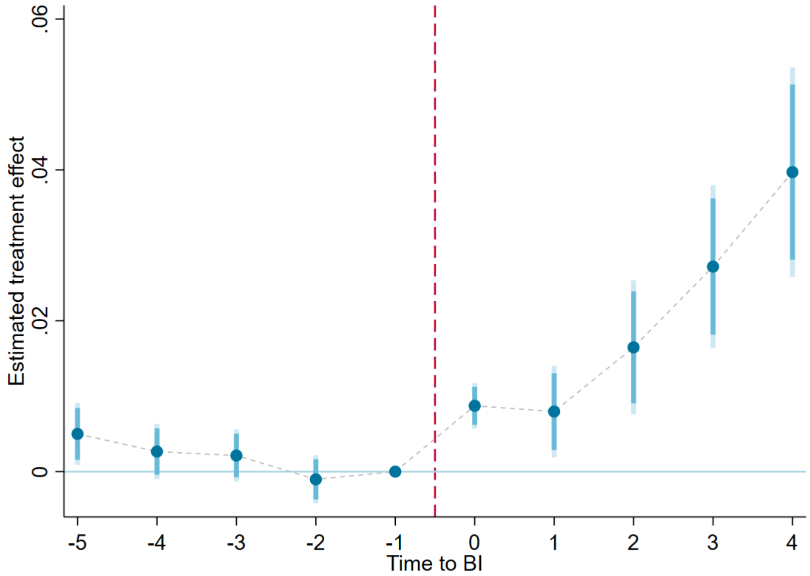

Notes: This Figure shows regression coefficients and 90% and 95% confidence intervals from a dynamic event study (see equation (3) and text for a description of the stacked event-study approach) where the dependent variable is the average ratio of the stock of capital in IT hardware over total capital by firms located in a given city in year t. Following definitions from the BIC-RN dataset, this ratio is defined as (LB-LC)/(BJ-BK) for each firm and then averaged at the city level. Pre-broadband expansion average of the dependent variable equals 7%.

4 Empirical Evidence on BI and Outsourcing

4.1 At the city and establishment level

We start by looking at the causal effect of BI on outsourcing expenditures and occupational segregation across firms. Figure A4 in Appendix A summarizes the evolution of these two outcomes over the period of interest. The outsourcing expenditures over wages bill grew very steeply at the end of the 1990s, going from 0.38 in 1997 to 0.44 in 2000 (a growth of 16%), then stabilized at this high plateau.[35] The level of occupation concentration in establishments increased steadily, from 0.33 in 1997 to 0.39 in 2007 (a growth of 18%). To test whether these trends are (partly) linked to the diffusion of BI, we run the empirical specification reported in Equation (3) on these outcomes. Figure

[35]Given that outsourcing can be a direct effect of productivity shocks, as described in our conceptual framework, we can imagine that the growth observed before the arrival of BI might be due to previous technologies, such as the diffusion of robots and computers. Additionally, we also observe a significant drop in prices in the services to firms sector during this period, which suggests that the volume of sales has actually increased over the period.

III presents the preferred event study graphs at the city level and at the establishment level (in blue), along with the robustness tests controlling for the propensity to adopt earlier interacted with year fixed effects (in red). Cities belonging to different cohorts of broadband expansion followed very similar trends before the arrival of the internet, but started spending more on outsourcing and became increasingly sorted after connection. This indicates that, after the arrival of BI, establishments within the city progressively specialized and employed fewer types of occupations in-house. At the same time, establishments increasingly bought these services from other firms. This has resulted in workers becoming increasingly segregated into firms that primarily hire their type of specialists. The establishment level results on outsourcing expenditures follow similar patterns but are smaller in magnitude. Hence, some of the effect measured may operate through composition: newly created firms spend more on outsourcing relative to disappearing firms. On the other hand, the establishment-level results on occupational segregation are of a similar magnitude to the city-level ones. Results remain unchanged when we control for the propensity of cities to adopt BI early, supporting the interpretation that the effect is driven by BI itself and not by differential trends explained by pre-existing characteristics.

Table A1 in Appendix A quantifies the effect at the city and establishment levels by presenting the dynamic post-BI coefficients.[36] Five years after the arrival of BI in the city, we observe an average increase in the outsourcing expenditures over wage bill of about 0.016, which corresponds to a growth of 6% with respect to the baseline levels, but it is marginally not significant. Occupational concentration increases by about 0.006 (1.6% growth relative to baseline) and the effect is highly significant. The effects at the establishment level are smaller but more precise on outsourcing expenditures: the latter increases by 0.014 (+4%), while occupation concentration increases by 0.007 (+1.9%) within existing establishments and the post-treatment coefficients

[36]The number of observations is smaller in the regressions on outsourcing expenditures because the sample is restricted to firms with more than 750,000 euros of turnover, and because the data only reports the headquarters location for multi-establishment firms.

FIGURE III. OUTSOURCING INTENSITY AND OCCUPATIONAL CONCENTRATION WITHIN ESTABLISHMENTS

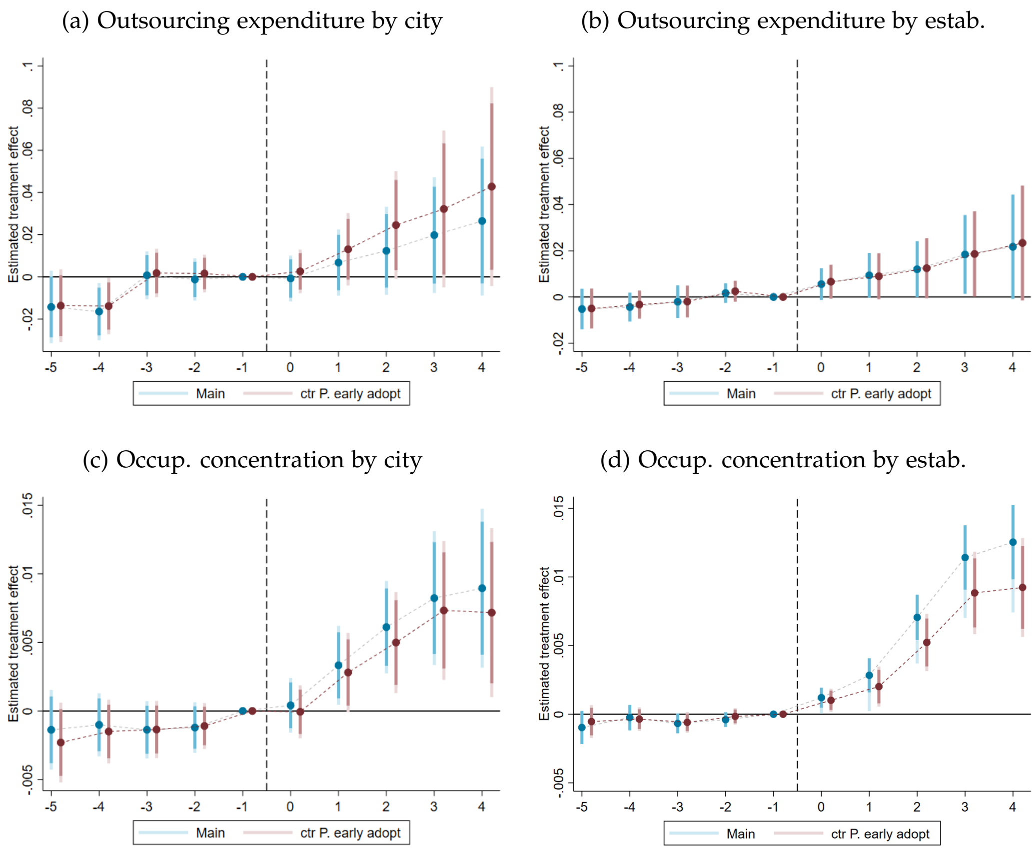

Notes: This figure shows regression coefficients and the 90% and 95% confidence intervals from a stacked event study design. The city-level specification follows Equation (3), while the establishment-level specification follows the same logic but replaces the city fixed effects with firm fixed effects. The blue lines present our baseline model, while the red lines present the model controlling for the propensity score of early adoption interacted with the year fixed effects.

remain strongly significant.

Next, we focus on the evolution of the business service sector. We consider two main measures at the city level, which are computed separately for the high- and low-skill segments: (i) the share of total employment in the city concentrated within outsourcing sectors, and (ii) the share of outsourceable workers in the city that are employed in outsourcing sectors.[37] Figure A5 in Appendix A describes the evolution of these outcomes over time. Employment in high-skill outsourcing services increased

[37] See Appendix B for a formal definition of the categories included in each group. High-skill outsourcing services include management, advertising, HR consulting, and IT services. Low-skill outsourcing services include security, cleaning, driving, and logistics. We match these sectors to their corresponding occupations.

by almost 50% over the years of BI diffusion, going from 6.2% of the labor force in 1997 to 9.2% in 2007. The share of employment in low-skill outsourcing services increased by 30% during the same period, going from 9% in 1997 to 11.5% in 2007. Beyond the overall growth of these two emblematic outsourcing sectors, we also observe a shift of outsourceable occupations towards these sectors and away from the rest of the economy. Specifically, the share of IT specialists and consultants employed by the high-skill service sector went from 17% in 1997 to 27% in 2007 (+59%). The share of security workers, cleaners, drivers, and logistic workers employed in the low-skill service sector went from 30% in 1997 to 53% in 2007 (+77%). In what follows, we test whether these substantial changes were, at least in part, facilitated by the diffusion of the ADSL technology across France.

Figure IV reports the event study graphs of the two main outcomes separately for high- and low-skill outsourcing, again showing the main model and the robustness to controlling for the probability of early adoption. Table A2 in Appendix A shows the corresponding regression coefficients, while Tables A3 and A4 show the coefficients obtained from breaking down each industry group even further into its sub-categories, underlining that similar patterns are observed for most of the sub-groups.[38]

Once again, the pre-trends are parallel and the effect materializes after the BI connection, regardless of whether we control for the propensity score. While the magnitude of the coefficients might appear small, the effect is actually far from trivial when compared to the baseline (pre-treatment) values.[39] The average share of city employment accounted for by high-skill outsourcing services was about 1.7% at the beginning of

[38]The regressions on the share of outsourceable occupations within outsourcing sectors have smaller and varying number of observations because they are only defined in cities where there is at least one firm active in outsourcing services, contrary to the regressions on the share of employment in outsourcing services, which is defined everywhere.

[39]Here we report the baseline averages computed across cities, which are our units of interest in this analysis. The magnitudes differ from the evidence presented in the summary graphs above because the latter report the values computed in the overall population (weighting larger cities more).

FIGURE IV. Effect of ADSL on high- and low-skill outsourcing

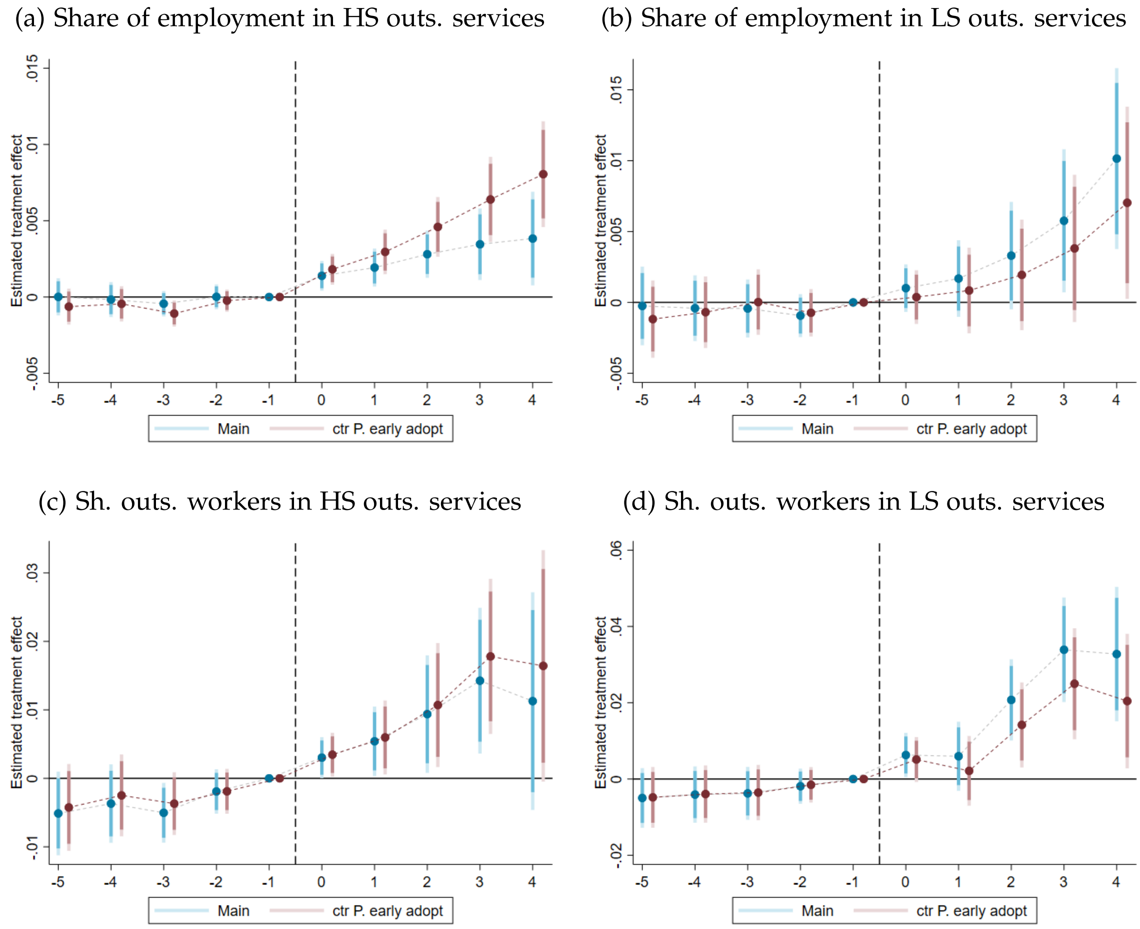

Notes: This figure shows regression coefficients and 90% and 95% confidence intervals from a stacked event study where the dependent variables are the share of workers employed in high-skill (low-skill) outsourcing services in a city at t (Figures 4(a) and 4(b)), and the share of outsourceable high-skill (low-skill) workers employed in their respective services in a city at t (Figures 4(c) and 4(d)). The model is presented in Equation (3). The blue lines present our baseline model, while the red lines present the model controlling for the propensity score of early adoption interacted with year fixed effects.

the period, such that the average impact of BI in the five years after its arrival amounts to a growth rate of 16%. Low-skill outsourcing services accounted for about 10% of total city employment in 1997, and the BI connection led to a growth rate of 4.6%. When it comes to the concentration of workers in outsourceable occupations within their respective service sectors, the baseline value is 6.5% for high-skill employees, and the effect of BI amounts to 1 additional percentage point (growth of 16%). For low-skill employees, the baseline value is 18% and the effect of BI is of 2 additional percentage points (growth of 12%). Given these magnitudes, we can infer that the arrival of the internet generated a structural change in the way these services are used by their business customers. This is an indication that BI catalyzed growth in domestic outsourcing, especially for non-core activities situated both at the low-skill and high-skill ends of the spectrum.

Robustness Tests Figures A7 and A8 in the Appendix show the robustness tests that we performed on the outsourcing outcomes at the city level, and Figures A10 and A11 in the Appendix show the same robustness tests on the outcomes at the firm level. In Figures A7 and A10 we test whether the results hold when adding additional controls: i) pre-BI productivity growth in the city interacted with year dummies, ii) sectoral composition in each city prior to 1999 interacted with year dummies, iii) share of left-wing voters in the city in 1995 and an indicator of cities switching the political majority between 1995 and 2002, both interacted with year dummies, iv) city size prior to 1999, measured as the number of establishments active in the city, interacted with year dummies, and v) all of the previous controls added together. The fact that the event study graphs remain highly similar for most outcomes suggests that differences in the timing of BI diffusion, across cities of similar density within the same department are not correlated with preexisting differences in city-level productivity growth, with differences in sectoral composition, with differences and changes in political affiliation, or with differences in city size. Figures A8 and A11 in the Appendix show the results obtained i) after controlling for the BI coverage of other cities in the same department - to get a sense of the possible spill-over effects of BI arrival on control cities; ii) after introducing commuting zone $\times$ year fixed effects instead of department $\times$ year, which should magnify the bias in the presence of strong geographic spill-overs; iii) from using the continuous measure of coverage $\tilde{Z}_{it}$ instead of the binary treatment based on the first year of BI arrival; and iv) from running a standard dynamic two-way fixed effects model, as reported in Equation (2). Once again, most of the outcomes remain unchanged, except that outsourcing expenditures becomes flat in the standard staggered event study model. This change suggests that dynamic effects in outsourcing expenditures bias downward the estimates in a model that uses post-treatment observations as controls.

Thirdly, Figures A9 and A12 present the results obtained from placebo tests, where late receivers are split into two groups according to their propensity of early adoption, and the pseudo-treatment is evaluated for the period preceding their actual BI connection. The graphs show flat and non-significant differences across the two groups except for one outcome: the share of employment in low-skill outsourcing services. This presents a positive trend among cities with higher likelihood to be treated early. We conclude that the causal impact of BI on this particular outcome necessitates a cautious interpretation. However, the lack of trends observed in the share of low-skill workers employed in low-skill services confirms that BI did generate a shift in distribution of outsourceable occupations across sectors, taking the overall size as given.

Fourthly, Table A7 reports the coefficients obtained from static regressions where the treatment status is interacted with the post-BI period, and Table A8 shows the coefficients from a standard staggered regression run at the commuting zone level, where the post-BI dummies are interacted with the continuous measure of BI coverage in the area. The magnitude of the static coefficients is slightly smaller than the average of the dynamic coefficients but remains in line with the main results. The effect measured at the employment zone level is smaller and less significant, which is explained by the fact that there is much less variation in the timing of BI arrival at this more aggregate level. Nonetheless, the coefficients qualitatively confirm those reported in the main analysis.

Finally, to further validate the stacked design, we estimate an alternative model using "pure controls" and we perform additional pre-trends tests as suggested by Borusyak et al. (2021).[40] For the first, we restrict the estimation period to only include years up to 2004 and then use cities treated after 2004 as pure controls, using both the stacked design and the standard TWFE design. Results are presented in Appendix figure

[40] We present these tests only for the city-level estimates for the sake of conciseness.

A13 and show very similar coefficients to our baseline estimates. For the second, we re-estimate both the dynamic stacked model and the dynamic TWFE model on the untreated sample - i.e. on the years in which a given city is not yet connected to BI -, to further validate the robustness of the common trend assumption. Because of multicollinearity, in the TWFE model we normalize both $t-4$ and $t-1$ to 0, while we only normalize $t-1$ to 0 in the stack model. Results are presented in Appendix figure A14. We cannot reject the fact that the pre-trend coefficients are jointly equal to 0.

Interpretation of effects To get a more concrete sense of the magnitude of these results, we can wonder how outsourcing would have evolved in the absence of BI.

To answer this question, we do some back-of-the-envelope calculations to compute the predicted trends in the main outcomes of interest after the subtraction of our estimated effects of BI. We follow the procedure adopted in Malgouyres et al. (2021) and we compute the predicted outcome as the actual outcome minus the dynamic effects predicted by our semi-dynamic specification reported in Table A1.

In this exercise, we partition cities into groups that receive broadband in the same year. We refer to these groups as cohorts and denote the first year of BI accessibility as t_0. The contribution of each cohort t_0 to the average effect at time t is the product of two terms. The first term is the weight of the cohort the year prior to its first connection (denoted $w_{t0,t0-1}$), which is computed as its share of total employment in that year. The second term is the average effect to which the cohort is exposed at time t. This effect will depend on the time since treatment $d = t - t_0$. We denote this effect as $\widehat{\alpha}_{t-t_0}$, which corresponds to the coefficient from our estimate of equation (3). We can then obtain the average effect of broadband internet expansion for each year t by summing across cohorts:

$$\overline{a}_t = \sum_{t_0=1999}^{2007} w_{t0,t0-1}\hat{\alpha}_{t-t_0}.$$

We postulate that the observed outcome y is given by a baseline level $y_t(0)$ that would have occurred in the absence of broadband diffusion, summed to the predicted effect of BI: $y_t = y_t(0) + \bar{a}_t$. We then obtain the trends in the main outcomes in the absence of our estimated BI effect by inverting this relationship: $y_t(0) = y_t - \bar{a}_t$. This exercise is not a proper counterfactual analysis since it ignores spill-overs and general equilibrium effects in the economy, which we suspect might be important in this case. However, we still believe that it can provide a useful benchmark to interpret the magnitude of our results.

Figure V shows that the outsourcing expenditures over wage bill decreases after the beginning of the 2000s if we subtract the estimated BI effect, and by the end of the period it is 9% smaller than what is actually observed. The occupational segregation across firms and the share of high- and low-skill outsourceable occupations in outsourcing sectors increase less steeply over the period when we subtract the estimated BI effect, reaching in 2007 levels that are 2.5%, 4.9%, and 6.6% smaller, respectively, than what is observed in the data.

4.2 At the individual level

In this section, we utilize our individual panel data to track workers over time and examine their job mobility choices. Specifically, we investigate the impact of BI expansion on the likelihood that workers change jobs. We take into account three types of job mobility: any establishment change, a transition from an establishment outside the outsourcing sector to one within the outsourcing sector, and a shift between two establishments outside the outsourcing sector.

In contrast to the previous model, workers are able to relocate from their current city over the years, which could result in them being treated multiple times. To prevent this, we assign a unique city to each individual i, representing the location where they first experience broadband internet access (during a year denoted as $t_{0,i}$). We also use this year to define the skill of the workers based on their occupation at $t_{0,i}$.

FIGURE V. COMPUTATION OF THE PREDICTED AGGREGATE TRENDS AFTER SUBTRACTING OUR ESTIMATED BI EFFECT

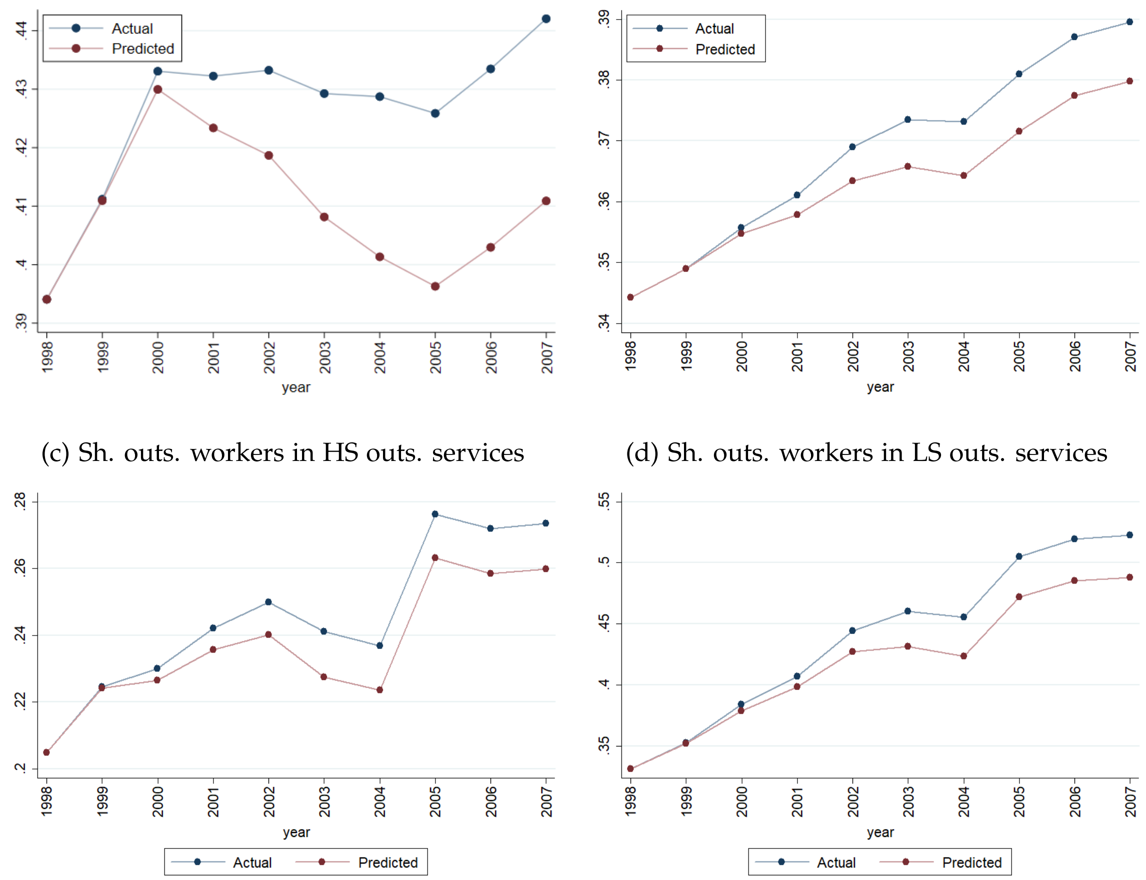

(a) Outsourcing expenditures

(b) Occupation concentration

(c) Sh. outs. workers in HS outs. services

(d) Sh. outs. workers in LS outs. services

Notes: The actual outcomes are the weighted averages of the outcomes observed in our sample aggregated at the economy level. The predicted outcomes are obtained by subtracting the predicted effect of broadband internet to the actual outcomes. The latter is computed using a weighted average of the estimated α_τ, for $\tau > 0$, where the weights correspond to the share in national employment for each cohort of firms (i.e., all firms for which broadband expansion occurred in the same year), measured in the year of broadband expansion. The weights are normalized so that they sum to one.

We then estimate the following linear probability model:

$$Move_{i,t} = \alpha \mathbb{1}\left\{t \geq t_{c(t_{0,i})}\right\} + \beta Z_{i,t} + \gamma_t \boldsymbol{X}_{\boldsymbol{c}(\boldsymbol{t}_{\boldsymbol{0},\boldsymbol{i}})} + \varepsilon_{i,t}, \tag{4}$$

where $c(t_{0,i})$ denotes the city in which the individual i worked at $t_{0,i}$ and $X_{c(t_{0,i})}$ is a vector of characteristics specific to the city that includes department fixed effects and the logarithm of density (both interacted with year fixed effects γ_t). Finally, $Z_{i,t}$ denotes individual characteristics: age, age squared, sex, and a dummy for working part-time (as opposed to full-time). The dependent variable $Move_{i,t}$ is a binary variable that equals 1 if the worker i has moved in year t and 0 otherwise. ε is an idiosyncratic error that we assume can be correlated within departments but not across them. α thus captures the effect of being connected to BI on the probability of moving, controlling for observable and time-unvarying unobservable worker characteristics. We augment this model by adding successively city fixed effects $c(t_0, i)$ and individual workers fixed effects. In these more saturated models, the identification is fully driven by the timing of BI expansion.

TABLE III. Effect of ADSL on workers' mobility

	High-skill workers			Low-skill workers		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Mobility	Any move	To outsourcing	Other	Any move	To outsourcing	Other
Panel A						
Connected	0.046***	0.006***	0.040***	0.052***	0.002***	0.050***
	(0.010)	(0.002)	(0.008)	(0.003)	(0.000)	(0.003)
Relative	**0.141**	**0.176**	**0.137**	**0.214**	**0.320**	**0.211**
Panel B						
Connected	0.009*	0.002*	0.007	0.023***	0.002***	0.021***
	(0.005)	(0.001)	(0.004)	(0.003)	(0.000)	(0.002)
Relative	**0.028**	**0.075**	**0.022**	**0.095**	**0.243**	**0.090**
Panel C						
Connected	0.011**	0.003**	0.008***	0.015***	0.001***	0.013***
	(0.005)	(0.001)	(0.004)	(0.003)	(0.000)	(0.002)
Relative	**0.034**	**0.080**	**0.029**	**0.060**	**0.177**	**0.056**
Average pre-treatment	0.325	0.032	0.293	0.242	0.007	0.235
Obs.	571,720	571,720	571,720	2,666,453	2,666,453	2,666,453

Notes: This Table presents the point estimate and standard errors of coefficient α in equation (4) as the ratio of α divided by the average value of the dependent variable pre-treatment (**Relative**). The dependent variable is a binary variable equal to 1 if worker i has changed establishment between t and $t+2$. Panel A includes department-year fixed effects, the logarithm of the density of the city interacted with a time dummy and individual controls: gender, age, age squared and a full-time dummy. Panel B adds a city fixed effect and panel C a city and individual fixed effects. Columns (1) and (4) consider any move, columns (2) and (5) consider mobility to a firm in the outsourcing sector and columns (3) and (6) consider mobility to firms excluding the outsourcing sectors. The outsourcing sectors respectively denote the high-skill outsourcing sectors in column (2) (e.g. IT services) and the low-skill outsourcing sectors in column (5) (e.g. cleaning services). Columns (1), (2) and (3) restrict to high-skilled occupation workers and columns (4), (5) and (6) to low-skilled occupation workers. OLS estimator with standard errors clustered at the department level and robust to heteroskedasticity. Time period: 1995-2008. ***, ** and * respectively denote significance at the 1, 5 and 10% level.

Table III presents our results. We report coefficient α from equation (4) for different types of mobility. Columns (1) to (3) refer only to workers in high-skilled occupations. In Column (1), the dependent variable is equal to 1 in the case of mobility, regardless of the sector of the destination firm. Column (2) is conditional on moving to a firm in any of the high-skill outsourcing sectors (IT, accounting, etc.) and Column (3) is conditional on moving to a firm outside these outsourcing sectors. Columns (4) to (6) do the same but for workers in low-skill occupations, and for the corresponding outsourcing sectors (cleaning, driving, security, etc.). Panel A does not include additional fixed effects on top of what is presented in Equation (4), Panel B includes city $c(t_0, i)$ fixed effects, and Panel C adds individual fixed effects. Because the average value of the dependent variable is very different across skills and types of mobility, in the table we also report the baseline average and the computed relative effect of BI in% terms (the line *Relative* in each panel), defined as the ratio of α over the pre-treatment average.

These results show that being connected to BI is associated with a greater propensity to move, both for high-skill and low-skill workers, with a larger relative effect for mobility to an outsourcing sector. Relatively to the pre-treatment average, and conditional on all covariates, being connected to BI induces a 0.141 percentage-point increase in the propensity to switch firms for workers in high-skill occupations, and a 0.214 percentage-point increase for workers in low-skill occupations. This relative change is larger for moves towards outsourcing firms (Columns 2 and 5), whether or not additional fixed effects are added.

Part of this mobility includes moves across cities, which could suggest that control cities in the specifications presented in Section 4.1 might be partly affected by the treatment through an increase in labor mobility. We argue that this effect is one order of magnitude smaller than the main outcomes of interest at the city level, and thus can be considered negligible. New entrant workers in low-skill occupations represent on average 0.5% of the overall employment of receiving cities and 1.3% of the outsourceable employment. New entrant workers in high-skill occupations represent on average 3.7% of the overall employment of receiving cities and 11% of outsourceable employment. In addition, the majority of the inter-city moves of both high-skill and low-skill workers are directed towards cities that are already connected to BI starting from 2001, and thus do not affect the controls (see Appendix Table A5).

Note that this model is different from the one used in Section 4.1 and therefore the identification discussed previously does not apply in the same way here. The advantage of using individual data is that it allows us to focus on the intensive margins of the effect and to look in more detail at the different types of mobility and how they are impacted by the arrival of BI. In Appendix Figure A6 we aggregate the data at the city level, using $c(t_0, i)$ for all individuals i, and we calculate the share of workers in high-skill and low-skill occupations that move to the outsourcing sector every year. We can then estimate the same model as in Section 4. To account for the fact that this model is aggregated from individual level data, we weight the city by population. These results support the causal interpretation of our individual-level regressions, namely that BI increases the mobility of workers from their current firms to firms specialized in outsourcing services.

5 Wage dynamics around individual transitions to outsourcing

Our findings suggest that BI expansion increases the mobility of workers across establishments, with a particularly strong effect on the mobility of workers in outsourceable occupations towards firms in the outsourcing sector. This reallocation of workers across firms might be another mechanism through which technological change affects the wages of low- and high-skill workers differently, on top of the skill-biased effect on the production function shown by Akerman et al. (2015) in Norway. In our context, we find broadly comparable skill-biased effects of broadband expansion on wages to those shown by Akerman et al. (2015)—results are presented in Appendix D. In this section, in order to assess more directly the relationship between mobility to outsourcing on wages, we carry out a series of simple event studies examining wage dynamics around *individual outsourcing events*, again distinguishing between high- and low-skill occupations. We define *individual* outsourcing events as follows:

- ***Individual outsourcing event.*** *The mobility of one worker from a firm outside the outsourcing sector to a firm belonging to the outsourcing sector.*

This definition allows us to capture many outsourcing events for both high- and low-skill workers. We thus look in greater detail at the wage dynamics of workers around these events, which helps indicate the working conditions in the outsourcing sectors compared to those elsewhere.

Table A6 in Appendix A describes the average wage and employment of outsourceable occupations across sectors. IT specialists and consultants earn a gross hourly wage of about 30€ in their outsourcing services, which is slightly higher than their counterparts in other services and in manufacturing (25€ and 28€, respectively). Security workers, cleaners, drivers, and logistic workers, on the other hand, earn less per hour when they are employed in their outsourcing services (12€) than when they are employed in other services or manufacturing (13€ and 15€, respectively). These characteristics are consistent with the idea that high-skill outsourcing may be voluntary, while low-skill outsourcing might not be.

To estimation the association between outsourcing and individual wages, we leverage our individual panel data and follow outsourced workers before and after an individual outsourcing event. We restrict our analysis to the subset of workers that experienced one and only one outsourcing event over the period of observation, which allows us to define our event study based on the timing of the move. We then look at the evolution of their hourly wages by estimating the following dynamic model:

$$\log(w_{it}^{ow}) = \sum_{\tau=-v}^{v'} \alpha_\tau \mathbb{1}\left\{t = t_i + \tau\right\} + X\gamma + \psi_{d,t} + \nu_i + \varepsilon_{i,t} \tag{5}$$

where w_{it}^{ow} represents the hourly wage of workers that are being outsourced by their firms, and t_i is the year of the event. X is a vector of time-varying individual characteristics: age, age squared, and an indicator of whether the job is part-time. $\psi_{d,t}$ and ν_i are a set of department d times year t fixed effects, and individual fixed effects.[41] ε is an idiosyncratic error that we assume can be correlated within departments but

[41] Adding sector $s(i)$ fixed effects does not affect our results.

not across them. We estimate this regression separately for low-skill workers and high-skill workers, with respective samples of 7,995 and 10,404 different workers. We consider $v = v' = 6$, which means that we follow workers over 13 years. Finally, given that in this setting, and similar to the model presented in Equation (2), every worker is treated, we standardize $\alpha_{-6} = \alpha_0 = 0$ (see Borusyak et al., 2021). We interpret the findings from this exercise as suggestive, as they rely on the relatively strong assumption that the timing of the outsourcing events is exogenous from the point of view of workers' wage trajectories, after controlling for observable characteristics and fixed effects. Still, the absence of pre-trends is indicative of the comparability of wages before workers' mobility to a service firm and accordingly, we believe the changes in wages around the events are interesting from a descriptive standpoint. Results are presented in Figures 6(a) and 6(b).

This figure shows that the hourly wage of outsourced low-skill workers sharply decreases after the outsourcing event. This finding is in line with the results documented by Goldschmidt and Schmieder (2017), who explain this phenomenon by the fact that firms in the outsourcing sectors benefit from lower rents, on average, than other companies. This translates into a lower wage premium for their employees.[42] The average decline ranges from -4% immediately after the outsourcing event to -2% five years later, gradually converging towards the pre-treatment level. By contrast, outsourced high-skill workers enjoy a 4% gain in hourly wages in the long run after the outsourcing transition, a pattern that is consistent with broadband stimulating demand for IT services and thus resulting in an increase in IT workers' outside options and voluntary job-to-job transitions. These particular high-skill occupations indeed continue to be in high demand and can easily be contracted, especially as communication technologies improve. By regrouping these occupations, specialized firms can serve

[42]The shaded area corresponds to the year of the move, when the outsourced worker quit his or her previous employer to join an outsourcing firm. This transition period is known to create reporting errors in the number of hours worked and the wages in the data, resulting in a possible noisy estimate for the dummy α_1.

FIGURE VI. Wages of outsourced workers before and after the outsourcing event

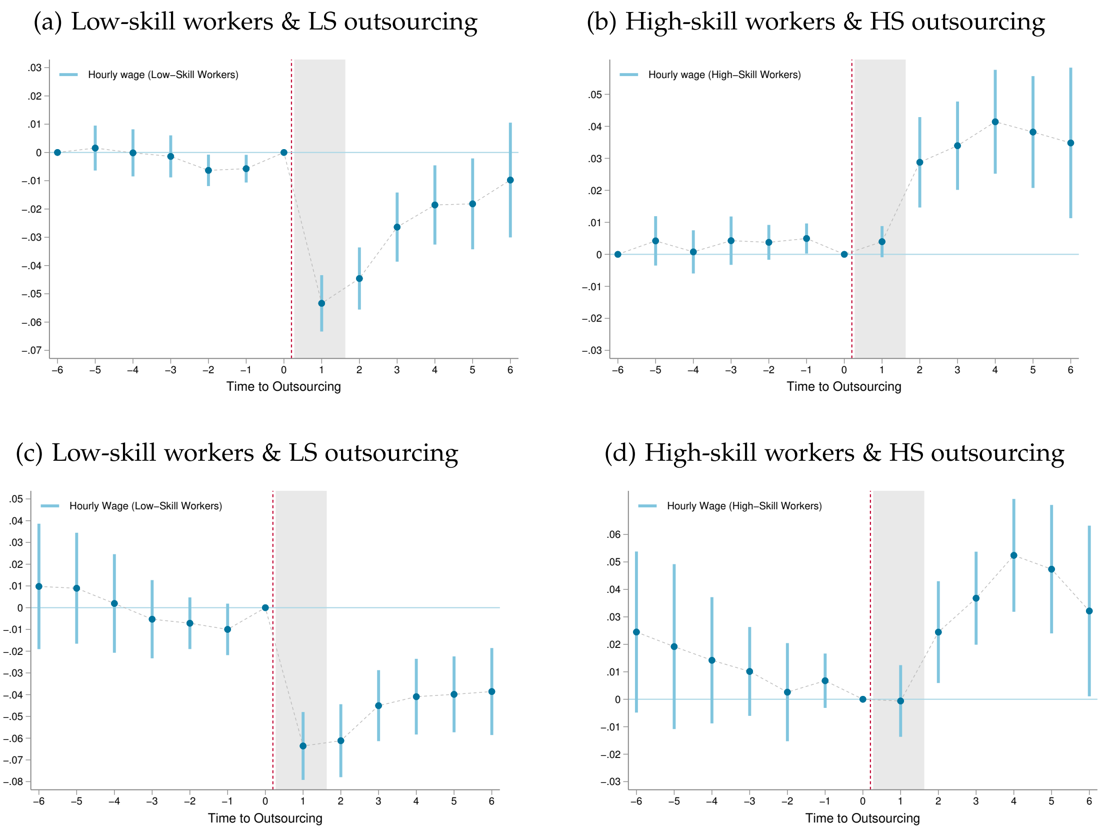

Notes: This figure shows regression coefficients and confidence intervals at 95% from a dynamic event study, where the dependent variables are the log hourly wage of workers being outsourced from the establishment at time t, and the regressors are dummies for the number of years before/after the establishment experiences an outsourcing event. It also includes several control variables: age, age squared, a part-time dummy, and a set of year times department and individual fixed effects. The shaded area denotes the year of the mobility, which is associated with noisy measures of work duration. The top panels are restricted to workers with one and only one instance of mobility from a non-outsourcing firm to an outsourcing firm. The bottom panels add control groups composed of workers who move from any firm to a non-outsourcing firm. The left-hand side panels are restricted to low-skill workers (39,659 workers in Panel (a) and 173,788 in Panel (c)), while the right-hand side is restricted to high-skill workers (10,404 workers in Panel (b) and 39,659 in Panel (d)). Standard errors are computed using an heteroskedastic robust variance covariance estimators allowing for autocorrelation at the department level.

different clients and generate more profit, first by maximizing the utilization rate of their inputs (mainly labor) and second by reducing fixed costs. This mechanism is similar in certain respects to some low-skill outsourcing, but in the case of IT specialists or consultants, the outsourced workers have greater bargaining power and can capture part of the rent by commanding a higher wage. Overall, these results emphasize a heterogeneous impact of outsourcing on workers' wages. While low-skill workers suffer from a significant wage loss, high-skill workers seem to benefit from outsourcing, as they are able to capture part of the increased profit made by firms as a result of this cost reduction.[43]

To disentangle the specific effect of moving towards the outsourcing sector relative to other types of mobility, in Figures 6(c) and 6(d) we reproduce the same exercise but using the set of workers who transition to a different firm that does not belong to one of the outsourcing sectors as a control group. In this case, the coefficients should be interpreted as the difference in hourly wages compared to workers of a similar skill level who move to another firm outside the outsourcing sector at the same point in time. In both cases, the results are consistent with the view that outsourced low-skill workers are negatively affected by an outsourcing event while high-skill outsourced workers are not, and that this effect goes beyond what could be observed in other types of moves.

6 Conclusion

The widespread diffusion of high-speed internet has fostered many changes in how businesses operate, some of which differ from those observed in previous waves of technological change. In this paper, we examine the role that broadband internet played in incentivizing French firms to outsource some non-core activities, both in the high-skill and low-skill segment, and study its effects on affected workers. We leverage the staggered roll-out of broadband internet across the French territory to adopt an event study design and estimate the causal effect associated with the diffusion of this new technology. This approach compares similarly dense cities within a given department, which gained access to BI at different times. Our results show that the internet is not only skill-biased but also increases the degree of occupational concentration within establishments, by pushing firms to outsource activities with lower degrees of complementarity in production. This phenomenon touches both low-skill

[43]Higher wages for high-skill workers could also indicate compensation for greater job security risks. Companies may regularly require cleaning or driving services, but only hire consultants as needed, leading to higher wages for the latter due to their inconsistent work hours.

and high-skill occupations. Finally, we provide evidence that high-skill workers experience wage gains through outsourcing, while low-skill workers experience wage losses. These findings are suggestive that the skill-biased impact of broadband internet technology is partly driven by its role in facilitating domestic outsourcing of both low- and high-skilled workers, which has important implications in terms of labor market inequality.

References

Abraham, Katharine G and Susan K Taylor, "Firms' use of outside contractors: Theory and evidence," *Journal of Labor Economics*, 1996, *14* (3), 394–424.

Abramovsky, Laura and Rachel Griffith, "Outsourcing and offshoring of business services: How important is ICT?," *Journal of the European Economic Association*, 2006, *4* (2-3), 594–601.

_ , _ , and Helen Miller, "Domestic effects of offshoring high-skilled jobs: Complementarities in knowledge production," *Review of International Economics*, 2017, *25* (1), 1–20.

Aghion, Philippe, Antonin Bergeaud, Richard Blundell, and Rachel Griffith, "The Innovation Premium to Soft Skills in Low Skill Occupations," Technical Report 14102, CEPR 2019.

_ , _ , Timo Boppart, Peter J Klenow, and Huiyu Li, "A theory of falling growth and rising rents," Technical Report w26448, National Bureau of Economic Research 2019.

_ , _ , _ , Peter J. Klenow, and Huiyu Li, "Good Rents versus Bad Rents: R&D Misallocation and Growth," 2023. mimeo Stanford.

Akerman, Anders, Edwin Leuven, and Magne Mogstad, "Information Frictions, Internet and the Relationship between Distance and Trade," Memorandum, Oslo University, Department of Economics February 2018.

_ , Ingvil Gaarder, and Magne Mogstad, "The skill complementarity of broadband internet," *Quarterly Journal of Economics*, 2015, *130* (4), 1781–1824.

Andrews, Dan, Giuseppe Nicoletti, and Christina Timiliotis, "Going digital: What determines technology diffusion among firms?," Technical Report 2018.

Arcep, "Etude sur les équipements et usages des PME et ETI," 2016. Arcep.

Baker, Andrew, David F Larcker, and Charles CY Wang, "How Much Should We Trust Staggered Difference-In-Differences Estimates?," *Available at SSRN 3794018*, 2021.

Bartel, Ann P, Saul Lach, and Nachum Sicherman, "Outsourcing and technological change," 2005.

Berlingieri, Giuseppe, "Outsourcing and the Rise in Services," Technical Report 1199, Centre for Economic Performance 2014.

Bernhardt, Annette, Rosemary Batt, Susan N Houseman, and Eileen Appelbaum, "Domestic outsourcing in the United States: a research agenda to assess trends and effects on job quality," Technical Report 16-253, Upjohn Institute Working Paper 2016.

Bertschek, Irene, Daniel Cerquera, and Gordon J Klein, "More bits–more bucks? Measuring the impact of broadband internet on firm performance," *Information Economics and Policy*, 2013, *25* (3), 190–203.

Bilal, Adrien and Hugo Lhuillier, "Outsourcing, Inequality and Aggregate Output," 2021. Mimeo Harvard.

Biscourp, Pierre and Francis Kramarz, "Employment, skill structure and international trade: Firm-level evidence for France," *Journal of International Economics*, 2007, *72* (1), 22–51.

Bloom, Nicholas, Luis Garicano, Raffaella Sadun, and John Van Reenen, "The distinct effects of information technology and communication technology on firm organization," *Management Science*, 2014, *60* (12), 2859–2885.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting event study designs: Robust and efficient estimation," 2021. Unpublished working paper, version May.

Callaway, Brantly and Pedro HC Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2020.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The effect of minimum wages on low-wage jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Chaisemartin, Clément De and Xavier d'Haultfoeuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–96.

Cortes, Guido Matias and Andrea Salvatori, "Delving into the demand side: Changes in workplace specialization and job polarization," *Labour economics*, 2019, *57*, 164–176.

Deming, David J and Kadeem Noray, "Earnings dynamics, changing job skills, and STEM careers," *The Quarterly Journal of Economics*, 2020, *135* (4), 1965–2005.

Deshpande, Manasi and Yue Li, "Who is screened out? application costs and the targeting of disability programs," *American Economic Journal: Economic Policy*, 2019, *11* (4), 213–48.

DeStefano, Timothy, Richard Kneller, and Jonathan Timmis, "Broadband infrastructure, ICT use and firm performance: Evidence for UK firms," *Journal of Economic Behavior & Organization*, 2018, *155*, 110–139.

Dorn, David, Johannes F Schmieder, and James R Spletzer, "Domestic Outsourcing of Labor Services in the United States: 1996-2015," 2018. Department of Labor.

Drenik, Andres, Simon Jäger, Pascuel Plotkin, and Benjamin Schoefer, "Paying outsourced labor: Direct evidence from linked temp agency-worker-client data," *The Review of Economics and Statistics*, 2020, pp. 1–28.

Dube, Arindrajit and Ethan Kaplan, "Does outsourcing reduce wages in the low-wage service occupations? Evidence from janitors and guards," *ILR Review*, 2010, *63* (2), 287–306.

Fabling, Richard and Arthur Grimes, "Picking up speed: Does ultrafast broadband increase firm productivity?," *Information Economics and Policy*, 2021, *57*, 100937.

Fort, Teresa C, "Technology and production fragmentation: Domestic versus foreign sourcing," *Review of Economic Studies*, 2017, *84* (2), 650–687.

Gardner, John, "Two-stage differences in differences," *arXiv preprint arXiv:2207.05943*, 2022.

Garicano, Luis and Esteban Rossi-Hansberg, "Organization and inequality in a knowledge economy," *Quarterly Journal of Economics*, 2006, *121* (4), 1383–1435.

Godechot, Olivier, Paula Apascaritei, István Boza, Lasse Folke Henriksen, Are Skeie Hermansen, Feng Hou, Naomi Kodama, Alena Křížková, Jiwook Jung, Marta M Elvira et al., "The great separation: Top earner segregation at work in high-income countries," Technical Report, MaxPo Discussion Paper 2020.

Goldschmidt, Deborah and Johannes F Schmieder, "The rise of domestic outsourcing and the evolution of the German wage structure," *Quarterly Journal of Economics*, 2017, *132* (3), 1165–1217.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," Technical Report, National Bureau of Economic Research 2018.

Grimes, Arthur, Cleo Ren, and Philip Stevens, "The need for speed: impacts of internet connectivity on firm productivity," *Journal of Productivity Analysis*, 2012, *37* (2), 187–201.

Haller, Stefanie A and Sean Lyons, "Broadband adoption and firm productivity: Evidence from Irish manufacturing firms," *Telecommunications Policy*, 2015, *39* (1), 1–13.

Hummels, David, Jakob R Munch, and Chong Xiang, "Offshoring and labor markets," *Journal of Economic Literature*, 2018, *56* (3), 981–1028.

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar, "Who profits from patents? rent-sharing at innovative firms," *Quarterly Journal of Economics*, 2019, *134* (3), 1343–1404.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler, "Imperfect competition, compensating differentials and rent sharing in the US labor market," Technical Report w25954, National Bureau of Economic Research 2019.

Lashkari, Danial, Arthur Bauer, and Jocelyn Boussard, "Information technology and returns to scale," *Available at SSRN 3458604*, 2022.

Lewis, Tracy R and David EM Sappington, "Technological Change and the Boundaries of the Firm," *American Economic Review*, 1991, pp. 887–900.

Malgouyres, Clément, Thierry Mayer, and Clément Mazet-Sonilhac, "Technology-induced trade shocks? Evidence from broadband expansion in France," *Journal of International Economics*, 2021, *133*, 103520.

Oi, Walter Y and Todd L Idson, "Firm size and wages," *Handbook of labor economics*, 1999, *3*, 2165–2214.

Ono, Yukako, "Outsourcing business services and the role of central administrative offices," *Journal of Urban Economics*, 2003, *53* (3), 377–395.

Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter, "Firming up inequality," *Quarterly Journal of Economics*, 2019, *134* (1), 1–50.

Spitz-Oener, Alexandra, "Technical change, job tasks, and rising educational demands: Looking outside the wage structure," *Journal of labor economics*, 2006, *24* (2), 235–270.

Tadelis, Steven, "Complexity, flexibility, and the make-or-buy decision," *American Economic Review*, 2002, *92* (2), 433–437.

Vannutelli, Silvia, "From Lapdogs to Watchdogs: Random Auditor Assignment and Municipal Fiscal Performance in Italy," *Job Market Paper*, 2020.

Weil, David, *The fissured workplace*, Harvard University Press, 2014.