



HAL
open science

Cutting hours through outsourcing

Olivier Godechot, Ulysse Lojkiné

► **To cite this version:**

| Olivier Godechot, Ulysse Lojkiné. Cutting hours through outsourcing. 2026. <halshs-05626506>

HAL Id: halshs-05626506

<https://shs.hal.science/halshs-05626506v1>

Preprint submitted on 19 May 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC 4.0 - Attribution - Non-commercial use - International License

CUTTING HOURS THROUGH OUTSOURCING

OLIVIER GODECHOT
ULYSSE LOJKINE

WORKING PAPER N°2026/10

APRIL 2026

WORLD
INEQUALITY
..... LAB

Cutting hours through outsourcing

Olivier Godechot & Ulysse Lojkine*

March 2026

Latest version [here](#)

Abstract

We study the on-site outsourcing of low-skilled service tasks (cleaning, catering and security) with French matched employer-employee data from the period 2001-2019. In line with the existing literature, we find a substantial penalty of around 10 log points in the yearly earnings of the outsourced employees, that persists 7 years after treatment. In contrast to the literature, we find that it is almost entirely explained by a penalty in days worked in the year and in hours worked per week, with at most a modest contribution of the hourly wage. We also find negative effects on employment, as proxied both by the presence in the panel and the probability of receiving unemployment benefits. Interviews with various stakeholders in these industries indicate that cutting the hours on a given site is a common way to compete on prices for subcontractor firms, with the decline in hours leading to both work intensification and a decrease in quality.

WIL working papers are circulated for discussion and comment purposes. Short sections of text may be quoted without explicit permission provided that full credit is given to the author(s). CC BY-NC-SA 4.0
How to cite this working paper: Godechot, O., Lojkine, U., Cutting hours through outsourcing, World Inequality Lab Working Paper 2026/10

*Sciences Po, CRIS-CNRS and AxPo. We are grateful to all the anonymised interviewees who lent us some of their time, to Cyprien Batut, François-Xavier Devetter, Daniele Girardi, Riddhi Kalsi, Julie Valentin and Michaël Zemmour for comments and suggestions, and to seminar participants at Sciences Po (Post-doc day 2025, CRIS Seminar, Labour reading group), PSE (Applied econ lunch seminar) and University Paris I (ADRES conference, EPCI seminar) for helpful discussion.

1 Introduction

There is increasing evidence that inequality between firms affects the labour market and wage distribution. Over the past decades, rising wage inequality between firms has contributed to an increase in total wage variance in Germany (Card et al., 2013) and the United States (Song et al., 2019), while in France, it has offset a compression within firms (Babet et al., 2026).

The same articles note a growing segregation of employees between firms, with lower-paid employees sharing less and less often the same employer with higher-paid employees (see Godechot et al., 2024). These trends can be partly explained by the increasing fragmentation of firms, or ‘fissuring’ (Weil, 2014), which gives rise to asymmetrical commercial relationships between firms. One form of it is temp agency labour, which, as Drenik et al. (2023) and Bergeaud et al. (2024) have shown on Argentinian and French data, partially excludes workers from the rent-sharing at play within companies.

Another important form, which is the subject of this paper, is the on-site outsourcing of low-skilled tasks, such as cleaning, catering and security. We examine that phenomenon in the French context. First, we use the Labour Force Survey (*Enquête Emploi*) to show descriptively that while these industries are small in total employment, they are important for some vulnerable segments of the workforce such as immigrant women without a degree; and that their employees tend to receive low earnings compared to other industries, due to both a low hourly wage and low hours, with an important share of them declaring that they would like to work more hours at the same hourly rate.

We then turn to our main event-study exercise. We use an improved version of Babet et al. (2026)’s procedure to chain administrative matched employer-employee data (*Déclarations Annuelles de Données Sociales*, DADS) into a panel for the period 2001-2019, which allows us to follow Goldschmidt and Schmieder (2017)’s definition of an outsourcing event as a group of workers moving at the same time from a common employer in some other industry to a new one in cleaning, catering or security.

With propensity score reweighting, we define a control group of non-treated or not-yet-treated employees who were similar in terms of observed demographic characteristics, worked in the same origin industry and occupation, and whose earnings and its components were similar in the four preceding years; we then estimate the dynamic treatment effect from this staggered treatment with the Local Projection Diff-in-Diff method from Dube et al. (2023).

We find a penalty of around 10 log points in yearly earnings for treated workers compared to the control group, that persists 7 years after treatment. In contrast to the findings of Goldschmidt and Schmieder (2017), that effect is not mostly driven by the hourly wage. We find an hourly

wage penalty starting two years after treatment, but it is small (1 to 2 log points), explaining only a minor part of the overall earnings effect, and it is significant only for some years and in some specifications.

The effect on earnings is almost entirely explained by a penalty in hours worked per year, both in terms of hours worked per day employed and of number of days worked. After treatment, the average gap between the two groups is of 1 to 1.5 hours per week on the former dimension, and of 3 to 10 days per year on the latter. Heterogeneity analyses reveal that, in the short term, the penalty is driven by workers who voluntarily or involuntarily left the destination employer after the transfer; however, in the medium term, stayers experience a penalty similar to the average. The analyses also show that the penalty is larger for women than for men, and for migrants than for natives.

The penalty in days worked per year can be interpreted as a negative effect on employment. It is complemented by two other metrics: a positive effect of 0.5 to 1 percentage point on the probability of receiving unemployment benefits but no earnings; and a negative effect on the probability of having any employment during the year. Again, these effects are stronger for women than for men, and for migrants than for natives.

The effects on hours and employment are, to the best of our knowledge, a novel finding in the economics literature on the effects of outsourcing. We suggest an interpretation based on 10 interviews conducted with managers, employees and union representatives at subcontractors in cleaning and catering. First, they uncover the role of the so-called “mobility clause”, that allows subcontracting employers to re-assign their employees across large geographical zones, shaping their bargaining power and allowing them to cut hours and jobs. Second, they show the importance of such site-level cuts in hours and jobs in the bidding process in which several subcontractors compete to obtain the contract from a client. Finally, they suggest that the consequences of these cuts are shared between employees (through work intensification) and the principal company or its customers (through a decrease in quality).

Literature

This paper stands at the intersection of two literatures. First, it contributes to the literature on the causes and effects of outsourcing. An early correlational study is that of Dube and Kaplan (2010) who find a wage penalty of 4% or more for outsourced janitors and 8% or more for outsourced guards. The seminal paper for an event study of low skilled service outsourcing is Goldschmidt and Schmieder (2017) on Germany, who find a 10 to 15 % daily wage penalty for workers outsourced into food, cleaning, security and logistics establishments. Applying a similar method on Turkish data,

Gürer and Taymaz (2025) find a stark contrast between low-skilled and high-skilled outsourcing: workers experience a penalty under the former, but a premium under the latter. Colonna and Aldeco Leo (2024) and Estefan et al. (2024) both study outsourcing in Mexico: they show that it was a way to avoid a within-company profit-sharing law, and that an outsourcing ban in 2021 increased compensation for insourced workers. In Brazil, Felix and Wong (2024) study the 1993 lift of an outsourcing ban and find substantial reallocation effects on security guards, where older in-house employees were fired and experienced persistent earnings losses while the task was contracted out to younger employees who benefited.

Several papers have studied outsourcing in France. Matching the REPOSE firm survey with DADS administrative wage data across all industries, Perraudin et al. (2014) show that the subcontractor status of the employer predicts lower wages for observationally similar employees; with the same data sources, Aeppli (2025) shows that it also predicts more employment instability. The causes of outsourcing are still understudied, although Bergeaud et al. (2025) show that when broadband internet reached a new city in France during the period 1999-2007, domestic outsourcing increased in that city, with negative effects on low-skilled and positive effects on high-skilled workers. Bilal and Lhuillier (2021) study service outsourcing in France, both low-skilled (security, cleaning, food, interim, general administrative services, call centres) and high-skilled (accounting, law or consulting services). Running a worker- and firm-fixed effect regression *à la* Abowd et al. (1999) on the (restricted) DADS panel between 1996 and 2007, they find that firm-level wage premia are 14 % lower at contractors compared to non-contractors. They also find that compared to non-contractors, contractors hire more from non-employment and experience a higher separation rate.

There is also a rich sociological and socio-economic literature. Weil (2014) provides a qualitative picture of the effects on workers of various forms of fissuring in the US, including outsourcing but also franchising. In France, outsourcing has been studied in the context of reception hostesses (Schütz, 2018), airport assistance (Brugière, 2017) and cleaning services (Puech, 2004; Devetter et al., 2021; Thevenot et al., 2021). In Canada, Zuberi (2013) studied the effects of the outsourcing of hospital cleaning on workers and on hygiene.

The second strand of literature to which we aim to contribute focuses on part-time work and its determinants. Studies show that part of the increase in German earnings inequality is explained by the increased prevalence of low hours at the bottom of the distribution (Checchi et al., 2016, 2022), while Beckmannshagen and Schröder (2022) use survey questions to show that much of it is involuntary part-time. Lachowska et al. (2023) analyse mobility patterns in the US and infer that many employees, especially at the bottom of the distribution, have a substantial willingness to pay part of their hourly wage for longer hours. Roux (2026) provides similar results on France,

complementing the more descriptive findings of Cohen et al. (2025). Other studies attempt to disentangle the sources of these constraints: while Borowczyk-Martins and Lalé (2019) show on US data that employers reduce workers’ hours during downturns to adjust to falling demand, Labanca and Pozzoli (2022) find evidence in Danish data of both constraints stemming from coordination between employees and from the firm’s “technology”. We contribute to this literature by showing how employers in outsourced service industries reduce workers’ hours to lower their clients’ costs.

Overview

In section 2, we provide descriptive statistics and the institutional context for low-skilled on-site service outsourcing in contemporary France. In section 3, we present our data and estimation strategy. In section 4, we turn to our estimates of the effects on earnings and its components and several proxies for unemployment, with special attention to the heterogeneity of these effects. Finally, in section 5, we interpret these results in the light of the existing qualitative evidence and the interviews we conducted.

2 Context and descriptive statistics

The rise of outsourcing across industries, and in particular in low-skilled services, has been documented in several advanced economies. In the US, between 1984 and 1999, Dube and Kaplan (2010) show an increase in the share of outsourced janitors from 16 % to 22 %, and of outsourced security guards from 40 % to 50 %. In Germany, Goldschmidt and Schmieder (2017) show a dramatic rise between 1975 and 2008, from less than 10% of cleaning and security jobs to almost 30% in security and almost 40% in cleaning.

In France, Bilal and Lhuillier (2021) show, based on the *Enquête Annuelle d’Entreprise* (a large annual firm survey) that outsourcing expenditures as a fraction of aggregate payroll increased from 6% in 1996 to over 10% in 2007, and, based on DADS, that the share of outsourced workers in occupations related to security, cleaning, food, interim, general administrative services and call centers rose by 16 pp, from less than 20 to more than 30% over the same period. Based on the *Enquête Emploi*, Devetter et al. (2021) show that the share of outsourced cleaners doubled from around 25 % in the early 1980s to more than 50 % in 2007, and fluctuating around that level ever since.

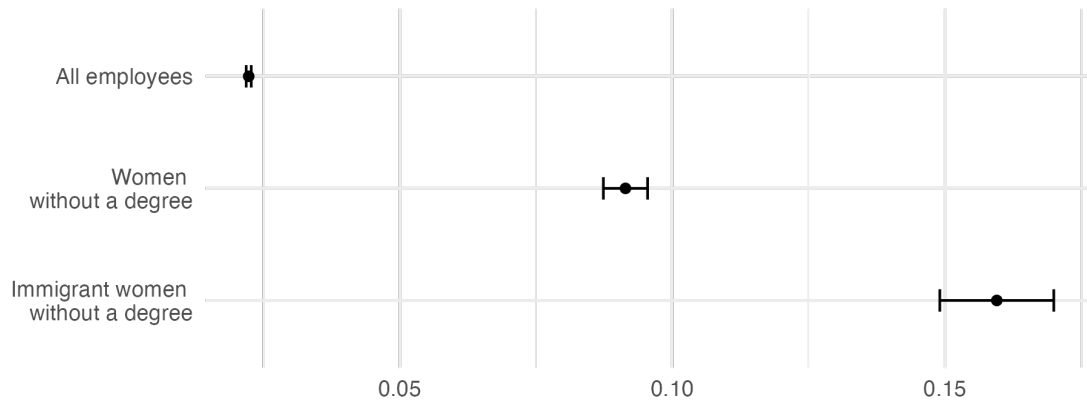


Figure 1: Share of the cleaning, catering and security industries among employees. Source: Enquête Emploi, 2013-2019.

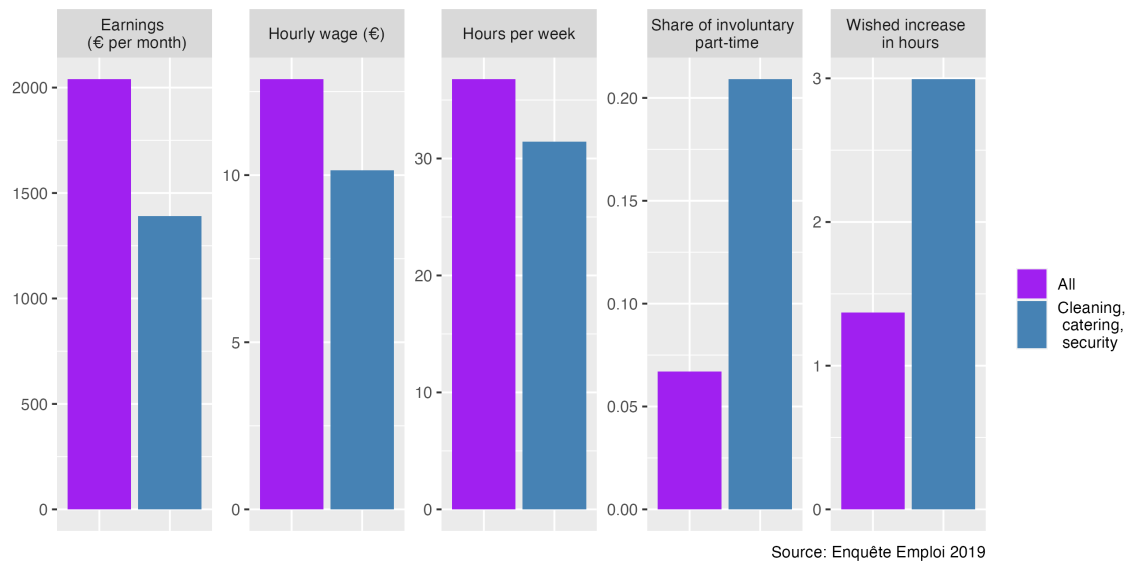


Figure 2

As shown by fig. 1, according to the French Labour Force Survey, these three industries together account for just 2 % of the employees, but that weight is much more important for certain

demographic groups, rising to 9 % among women without any degree (despite the overwhelmingly male workforce in security), and even to 16 % for immigrant women without any degree.

As shown in figure 2, according to the 2019 edition of the French Labour Force Survey, the earnings (here, the net earnings from their main job) of employees in the cleaning, catering and security industries are low — at 1390 euros per month, which is two-thirds of the average French employee. That gap is explained by both a lower hourly wage and fewer hours worked per week: 31.4 hours on average in these industries, substantially below the French legal working week of 35 hours, compared to 36.8 hours for the average French employee. The survey allows us to explore this last point further: 21% of respondents employed in these industries declare that the part-time is not voluntarily, but because their employer requires them to; if everyone worked their desired number of hours at their current hourly rate, the average working week in these industries would be 34.4 hours, a 10% increase.

This point can be illustrated by quoting the interview of Malika (see Appendix C), single mother of three, who earns 1000 to 1200 euros by working 22 hours per week, shared among two cleaning employers. A recent increase of about 10% in her hourly wage was “not much”, “not a change”, as contrasted with the months where she gets called for an extra 7-hour shift on sundays – then “I get more hours, I am lucky, I get 1500”.

Of course, a substantial part of these gaps can be explained by the characteristics of these workers, who are among the most vulnerable in the French labour market. 74% of the employees in those industries are not high school graduates (vs. 37% in other industries), 61 % are women (vs 50 %), and 30 % are immigrants (vs. 10%). All these factors contribute to limiting the options available for them on the labour market, as illustrated by our interviews in the HR team of a catering subcontractor. Claire, in charge of recruitment, highlights the ease of filling a food service assistant position (*employées de restauration*, jobs that do not require any formal qualification and not necessarily language skills): they receive “up to 600 or 700 applications” and never less than 10, so that the vacancy is “sometimes filled within a week”. She contrasts it with the difficulty in recruiting for other working-class occupations that require a diploma –, such as chefs, who need a vocational qualification (*CAP, Certificat d’Aptitude Professionnelle*), or, in the industrial company where Claire used to work before, skilled blue-collar workers who needed another, slightly higher qualification (*bac professionnel*). Her colleague Marc-André, in charge of training, stressed the limited French language skills of many food service assistants as restricting their career opportunities, particularly because higher positions in the kitchen require the ability to read instructions in order to apply hygiene guidelines.

A common feature of these three industries is the role played by calls for tenders, where the

principal periodically puts subcontracting companies in competition with one another to obtain the service at the best price. This leads to frequent substitutions of one subcontractor for another, governed by a provision common to the three industry-wide collective agreements¹: when the contract for a given site is awarded to a new provider, the employment contract of the outsourced workers on that site must also be transferred. We will come back to that feature in the interpretation of our quantitative results in section 5.

3 Method

3.1 Data

The original data source is BTS (*Base Tous Salariés* or Database of All Employees, formerly known as DADS, *Déclaration Annuelle de Données Sociales* or Annual Report of Social Data), i.e. wage declaration of all employment spells in France to Social Security. The Insee releases a “panelised” version for 8% of the sample only. But each year’s data also contains information on the previous year, which allowed Babet et al. (2026) to chain them and reconstruct a quasi-exhaustive panel for the period 2002-2019. The matching between annual files is based on establishment, gender, number of hours, start and end dates of the job, place of work and residence, earnings, and age. As explained in more detail in Appendix A.2, we build on their method while improving it on the margin by refining the matching on age and taking unemployment benefits into account. The matching rate on the enlarged domain is 97% on average between 2002 and 2019.

We construct:

- Total earnings in the year as the sum of the earnings of each employment spell (`S_BRUT`) of the individual in the year. These are gross or “posted” earnings, net of employer but not of employee Social Security Contributions. We express them in constant 2015 euros, using the Consumer Price Index from Insee.
- Days worked in the year by aggregating the durations of employment spells, taking overlaps into account (cf. Appendix A.1)
- Daily working hours as the ratio of total hours worked (sum of hours worked in each employment spell `NBHEUR`) to number of days worked. (Note that weekends included within a continuous employment spell are counted as days worked; multiplying by 7 converts this variable into hours per week, as in Figure 9.)

¹Cleaning: IDCC 3043, article 7. Catering: IDCC 1266, Avenant n° 3 du 26 février 1986. Security: IDCC 1351, Accord du 5 mars 2002 relatif à la reprise du personnel.

- Hourly wage as the ratio of total earnings to total hours.
- An employment dummy, which is equal to 1 when the individual has earnings in that year. In other words, an individual is considered as non-employed if she either appears in the panel solely as an unemployment benefit recipient, or does not appear at all.
- A stable establishment dummy. At time $t + h$, an individual is considered to be in the same establishment (as right after treatment) if her main establishment on Dec. 31st is the same as on Dec. 31st $t + 1$. At time $t - h$, an individual is considered to be in the same establishment (as right before treatment) if her main establishment on Jan. 1st is the same as on Jan. 1st t .
- A dummy for receiving unemployment benefits. Note that this information is reported in the data only between 2002 and 2015. Furthermore, unemployment benefits are reported only for individuals who would otherwise be present in the database, i.e. those with some paid employment the same year or the previous year.
- A dummy for receiving *only* unemployment benefits, i.e. receiving unemployment benefits but no earnings in the year.

3.2 Definition of the treatment and the control group

We identify an outsourcing event as a group of at least six persons who are employed on 1 January in the same non-FCSL establishment and on 31 December of the next year in the same catering, cleaning, or security establishment. The outsourced workers need to represent strictly less than half of the original establishment’s workforce, and the original establishment should still exist the next year. That destination establishment should not belong to the same firm or the same business group as the original establishment. We restrict ourselves to workers employed in the private sector the year before treatment².

With this definition, we identify 764 outsourcing events happening between 2005 and 2019 covering 4,558 treated workers³. Most of them were outsourced into catering or cleaning, and only a few hundreds outsourced into security. These events were spread over our whole period of study (2005-2019), although irregularly, with a maximum of 119 events in 2013 and a minimum of 25 events in 2009. Hospitals (by construction, private ones) are the most frequent industry of origin, common to 39 % of treated employees.

²The main establishment in the year before treatment should be either “Organismes privés spécialisés et groupements de droit privé” or “Autres sociétés privées”. This excludes public establishments, “entreprises individuelles” and “particulier employeur”.

³Because, in the main specification, the control group is built to be similar in terms of employment for $t - 1, \dots, t - 4$, we retain only events happening after 2005.

	To catering	To cleaning	To security	From a hospital	2005-2009	2010-2014	2015-2019	Total
# events	242	324	198	169	231	330	203	764
# individuals	1,996	1,956	606	1,782	1,705	1,633	1,220	4,558

Table 1: The treated group and its composition

To define an adequate control group, we first select a random subsample of 0.5% of all individuals present at least once in the panel in the period 2001-2019; after removing all observations of an already treated individuals (clean control condition) and imposing the presence of the individual at the beginning of year t and end of year $t + 1$ for comparability with the treated group, we obtain a control sample of 763 thousand observations. We then use propensity score weighting (also called inverse probability weighting, cf. Abadie (2005) and Appendix B below), i.e. we regress a treatment dummy on a set of predictors (controls), from which we estimate for each individual a predicted probability of being treated, and then re-weight individuals in the control group, assigning a larger weight to those with a larger predicted probability of treatment, i.e. more similar to the treated.

In our main specification, we control for, on the one hand, the range of variables at $t - 1$ – sex interacted with age and age squared, migrant status, establishment and firm size, seniority in establishment (for the main job), number of jobs in the year, 3-digit industry, 4-digit occupation –, and on the other hand, the following variables at each date between $t - 1$ and $t - 4$ included: employment status, log yearly earnings, log days and log hours per day. For reasons of comparability, we impose that members of the control group should work in a private establishment at $t - 1$ and, because that criterion is indirectly used in the definition of the treatment, that they should be employed on 1 January of t and on 31 December of $t + 1$.

1,030 individuals contribute to 10 % of the resulting reweighted control group and 11,343 individuals to 50 % of it⁴. Some characteristics of the treated and the control group are presented in table 2.

3.3 Estimation of the treatment effect

To estimate a treatment effect while avoiding the issues with the naive two-way fixed-effect regression applied in a staggered treatment setting, we combine the propensity-score weighting described

⁴To speak of “the” control group is a slight abuse of language. Here we are referring to the controls at horizon $h = 0$, but as detailed in Appendix B, a control group is defined for each horizon h : when moving from h to $h' > h$, the control group changes both because some individuals who were in the control group at h get treated and thus cannot serve as “clean” controls any more, and because some individuals who were in the treated group at h disappear from the panel at h' , changing the composition of the treated group and thus requiring new weights for the controls.

	Treated	Control
Age	44	44
Female	66 %	67 %
Migrant	24 %	24 %
Est. size	661	680
Firm size	3,087	3,275
Full time	55 %	55 %

Table 2: Average characteristics of the treated and control group the year before treatment

above with the LP-DiD method (Dube et al., 2023), a flexible and computationally light variant of Callaway and Sant’Anna (2021).

The equation estimated for each outcome y (log earnings, log days, log hours per day, log hourly wage, employment dummy, unemployment benefits dummy and stable establishment dummy) and horizon $h \in \{-6, \dots, 7\}$ is as follows:

$$y_{i,t+h} - y_{i,t-1} = \beta_h \Delta T_{i,t} + \delta_t^h + u_{i,h} \quad (1)$$

where y_i is the outcome considered, $\Delta T_{i,t}$ is an indicator for newly treated units at time t , δ_t^h is a year fixed-effect and $u_{i,h}$ is a random error term. All standard errors are clustered at the level of the firm identifier of the main job at $t - 1$.

As required by the method (Dube et al., 2023, p. 6), for β_h to yield an estimator of the equally-weighted average treatment effect on the treated (ATT), equation 1 is estimated each time with “clean controls”, i.e. removing from the sample the “forbidden comparisons” which are at the root of the distortions of the TWFE estimation in a staggered treatment context, and observations are, in addition to the propensity score weights mentioned above, re-weighted by a function of the proportion of treated units in each cohort (see Appendix B).

3.4 Complementary interviews

To complement our quantitative analysis, we conducted nine face-to-face interviews with stakeholders in the cleaning and catering industries, including managers, employees, and union representatives (see Appendix C for details). Initial interviewees were contacted through personal networks, and additional participants were recruited through snowball sampling. While this small convenience sample is not statistically representative, we aimed to include at least one representative from each key functional role (including workers, supervisors, employers, and union delegates) and sector (cleaning and catering). This approach allowed us to cross-reference experiences and perspectives across functional, sectoral, and class divides. The qualitative insights help illuminate

the mechanisms behind our quantitative findings, particularly how cost-cutting strategies shape workers' outcomes. We thus follow both the tradition of mixed-methods research in the social sciences (Pearce, 2012) and recent trends in economics, where qualitative insights help disentangle the mechanisms underlying causal estimates (Bergman et al., 2024).

4 Results

4.1 Average effects

Earnings and its components As shown by figure 3, we find a large and persistent effect on total yearly earnings of close to -10 log points at $t + 2$, reaching -13 log points at $t + 6$. The same figure shows that this earnings penalty is mainly explained by negative effects on the number of hours worked per day and days worked in the year. The former is more important in the short term and the latter in the longer term, although both effects are present and significant over the whole post-treatment window, with only one exception (the effect on hours per day ceases to be significant at $t + 7$). In contrast, the hourly wage explains at most a minor part of the effect on total earnings. Indeed, as shown in figure 17, we measure a significant effect on the hourly wage only for two years, $t + 5$ and $t + 6$, and even then, its magnitude of around -2 log points represents only a minor part of the total effect.

As detailed in section 4.3, that qualitative picture of a substantial and persistent fall in yearly earnings, with at most a small contribution of the hourly wage, is robust to several alternative specifications.

As discussed in section 5.4, that result stands in contrast to the seminal outsourcing event study from Goldschmidt and Schmieder (2017) on German data, which is all the more surprising since we borrowed from them the method for identifying outsourcing transfers in panel earnings data. Indeed, they measure a penalty in the daily wage and, although their data do not allow them to decompose it as we do, they give plausible reasons to interpret it as driven by the hourly wage.

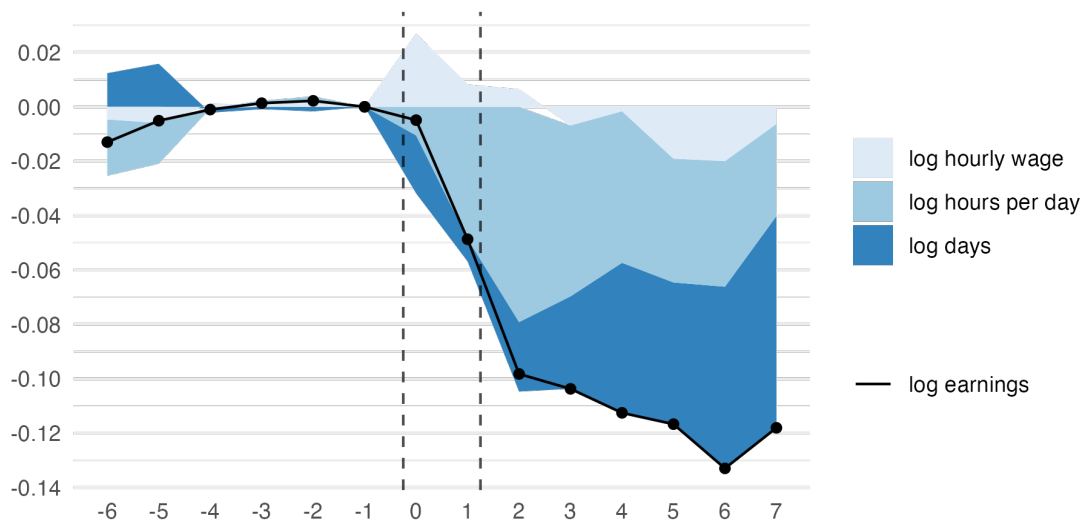


Figure 3: Effect of outsourcing on log earnings and its components. Main specification.




Figure 4: Effect on the probability of employment (sample restricted to those aged 50 or less at time of treatment) and of receiving unemployment benefits (sample restricted to 2002-2015)

Employment effects The penalty in the number of days worked in the year that we just mentioned may reflect both leaves of absence (parental leave and, perhaps more importantly here, sick

leave) as well as additional periods without employment.

We complement it by measuring the effect of treatment on two other outcomes: the probability of having any employment during a calendar year, and the probability of receiving unemployment benefits but no earnings during a year.

These measures have limitations. The main limitation of the employment dummy is related to the very nature of the database used, where an individual appears in year t only if she has a job in year t or had one in year $t - 1$. Because of the way we construct the panel, as detailed in Appendix A, that implies that we lose track of an individual if she is out of employment for a full calendar year and does not receive unemployment benefits during that year, or if she is out of employment for two consecutive calendar years, even if she receives unemployment benefits during these years. Because all units in the sample are employed at t and $t + 1$ by construction, that means that the employment dummy variable is reliable at $t + 2$ but relatively less reliable at longer time horizons, when individuals coming back into employment after one or more years are missed. Note two other limitations in terms of sample: to remove retirement effects, when measuring the effect on employment, we restrict the sample to individuals aged 50 or less at time of treatment; the unemployment benefits are missing from our data after 2015.

With these caveats in mind, the effects reported in Figure 4 are strongly suggestive of negative effects on employment. We find a significant negative and growing effect on the employment dummy, at around -2 pp in the short term and more than -5 pp at $t + 6$ and $t + 7$. We also find a positive effect of around $+1$ pp on the probability of receiving unemployment benefits, significant at $t + 3$, $t + 4$ and $t + 6$. The latter effect is small in the absolute, but high relative to the control group: as shown in figure 14, in some post-treatment years the share of individuals receiving unemployment benefits is actually twice as high in the treated as in the control group⁵.

A further finding deserves mention here: as shown in figure 19, we find a massive negative effect on the probability of keeping the same employer among those who remain in employment, of -20 pp at $t + 3$ and even larger in the longer term. This effect is, however, difficult to interpret, because we cannot isolate the contribution of each of the two underlying mechanisms that it may reflect. One is worker separations, i.e. employees voluntarily or involuntarily leaving their position after being outsourced and subsequently securing employment elsewhere, with the intervening spell of non-employment being consistent with the reduction in annual days worked documented above.

⁵As shown in figure 18, we also compute the effect on the probability of receiving unemployment benefits (including for those with earnings in the same year), but an important pre-trend makes the coefficients uninterpretable. However, when restricting the analysis to either catering or cleaning (fig. 30), the pre-trend disappears and a positive and significant effect remains, between $+2$ and $+5$ pp in catering in years $t + 2$ to $t + 7$, and between $+5$ and $+7$ pp in cleaning in years $t + 4$ to $t + 6$.

The second channel involves employer transitions without job transitions, when the worker’s employment contract is reassigned to an incoming service provider upon contract renewal — a common occurrence in subcontracting markets, as confirmed in our interviews — such that the employee changes employer while remaining in the same role.

4.2 Heterogeneity analysis

Balanced panel Since the share of individuals remaining employed after treatment declines rapidly following treatment (see Figure 12), the effect on earnings at h is estimated on the subsample of treated individuals still employed at that date. This approach makes full use of the available information, but it precludes interpreting the sequence of average effects as the effects experienced by the average treated individual. For this reason, we also estimate, as is common practice, effects within a balanced panel — specifically, a dynamically balanced panel, i.e. retaining only individuals who are continuously observed between $t - 3$ and $t + 5$ (we narrow the time window relative to the main specification to avoid excessive loss of statistical power).

The results are shown in Appendix F.1. We find a significant effect on earnings as early as $t + 2$, but much smaller than in the main specification. It grows over time and exceeds -10 log points by $t + 5$. As in the main specification, this total effect is driven by days worked and hours per day, rather than by the hourly wage, on which we find no significant effect in this specification.

This implies that in the treated group, the most severely affected in the short run tend to subsequently drop out of the panel, with the employment effect thus concentrating — consistent with intuition — among individuals who had already been negatively affected in terms of wages. The remaining individuals, those who remain employed later, experience more moderate short-run effects, but their situation (relative to the control group) then deteriorates over time.

Stayers As noted above, and as can be seen in Figure 15, the share of individuals who remain at the same establishment declines over time, particularly so in the treated group. This means that the average effect on earnings combines the effect on individuals who have stayed employed at the same subcontractor after having been transferred there at the time of treatment, and the effect on individuals who have voluntarily or involuntarily left that job for another.

This second mechanism is of particular interest, as it reveals the practices of subcontractors toward the workforce they take on following a transfer. To isolate it, we run the analysis on a sample that, at each horizon $h \geq 1$, is restricted to individuals who have the same employer at $t + h$ as immediately after treatment (and at $h \leq 0$, to individuals who have the same employer as immediately before treatment; see the description of the stable establishment dummy above in

Section 3.1). Note that this is a conservative definition of stayers: we retain only those who keep the same employer, but not – as they cannot be identified – those who remain at the same site under a different service provider following a contract renewal.

The results are presented in Appendix F.2. In the short run, as in the balanced panel, we find significant negative effects on earnings, but more moderate than in the main specification, indicating that the individuals most affected in the short run are those who immediately left the establishment to which they were transferred. From $t + 4$ onward, however, the effect reaches -10 log points: in the medium run, having been willing or able to remain at the destination establishment of the transfer does not provide protection against a penalty in earnings of roughly the same magnitude as the average. Once again, that penalty is driven by hours worked in the year.

By destination industry We run the analysis separately for each of the three destination industries: catering, cleaning, and security. Results are presented in Appendix F.3.

As mentioned above (Table 1), we identified only 198 outsourcing events towards a security subcontractor, representing 606 individuals; it is therefore probably unsurprising that we do not find any significant effect on this subsample. The comparison between catering and cleaning is more informative. The effects on earnings, on the employment dummy and on the probability of receiving unemployment benefits are broadly similar between catering and cleaning, although the earnings penalty is consistently higher in cleaning. The most striking difference is in the composition of that effect. We find a substantial and significant hourly wage penalty for outsourced cleaners, that explains an important part of the total earnings loss. By contrast, we do not find any such penalty in catering, and we even find a positive, significant effect of $+2$ to $+3$ log points between $t + 2$ and $t + 4$. This means that the null to small negative effect on the hourly wage presented above for the full sample is actually the average between a positive effect in catering and a negative effect in cleaning.

To conclude this paragraph, it is worth mentioning the outsourcing of logistics. Because it is included in the seminal Goldschmidt and Schmieder (2017) paper, we also identified these events and ran a separate analysis on that subsample. Results are in stark contrast to those of our main sample: we do not find any negative effect on earnings, on the employment dummy or on the establishment stability dummy, neither any positive effect on the probability of receiving unemployment benefits. This contrast is maybe not so surprising, given the different characteristics of the industries. Catering, cleaning and security have in common that they involve on-site service provision, which is not always the case in logistics (separate warehouses, transportation); they involve regular renewal of subcontracting contracts, upon which workers are transferred from one subcontractor to another, as stipulated in the collective agreement of each of these three sectors — a mechanism

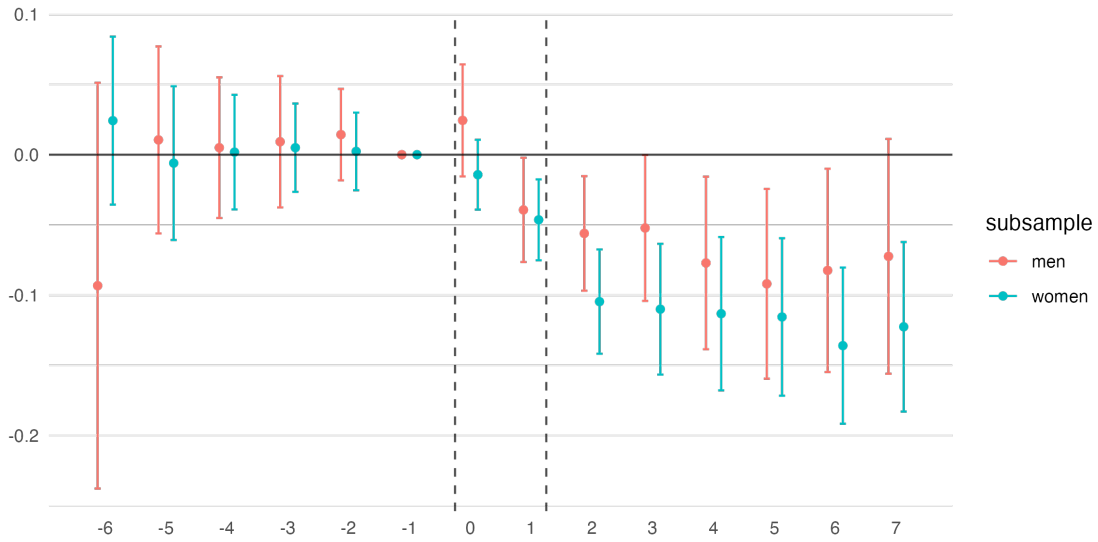


Figure 5: Heterogeneity by sex. Log earnings

that arguably generates a specific competitive pressure on subcontractors and therefore on workers, which may be less the case in logistics. Still, the result stands in contrast to that of Goldschmidt and Schmiuder (2017), who found a significant and substantial (although milder than in cleaning, see their table II) earnings penalty for outsourced workers in logistics.

By sex As shown in Figure 5, the earnings penalty is consistently larger for women (fluctuating between -10 and -14 log points) than for men (fluctuating between -6 and -9 log points). As shown in Appendix F.4, this is explained by a larger effect both on days and, except at $t + 4$ and $t + 5$, on hours per day.

The contrast on the employment penalty is even more striking. While we find large employment effects on women, ranging from -3 pp at $t + 2$ to -9 pp at $t + 6$, we do not measure any effect on the employment dummy of men. While we accordingly find higher effects on women’s probability of receiving unemployment benefits between $t + 3$ and $t + 4$, this ceases to be true at later dates. Among possible explanations are that these women, although seeking a job, could be less inclined than men to take up their unemployment insurance, or that they may be more likely to leave the labour force.

By migration status As shown in Figure 6, the earnings penalty is consistently larger for migrants than for natives. This is especially true in the medium term, with a massive penalty of

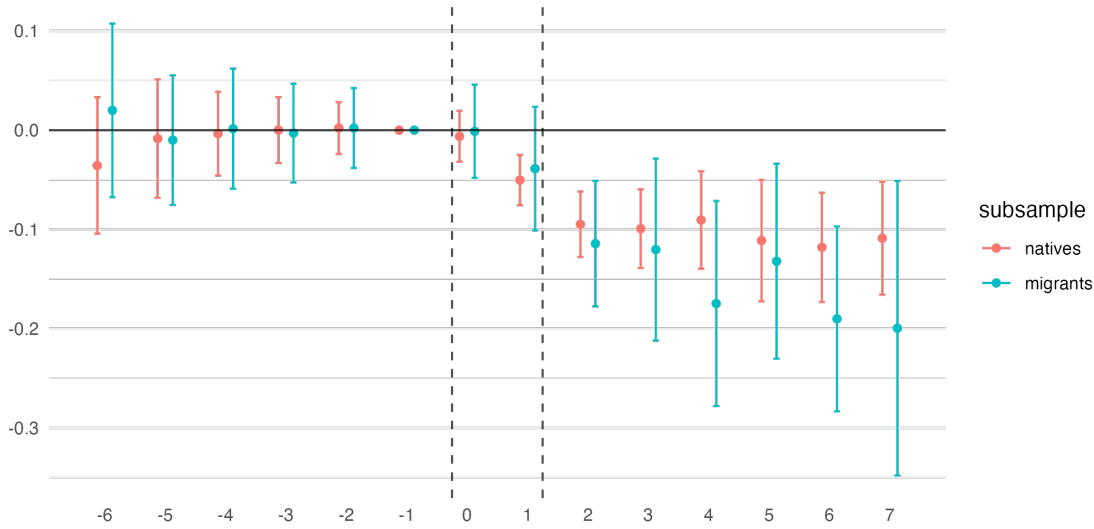


Figure 6: Heterogeneity by migration status. Log earnings

-17 to -19 log points for migrants at $t + 4$, $t + 6$ and $t + 7$, compared to -9 to -12 log points for natives at the same horizons. As shown in Appendix F.5 gap is explained mostly by a more important effect on hours per day and on the hourly wage (along with cleaners, migrants represent another subsample where we measure a significant hourly wage penalty, of -4 log points in the medium term).

Migrants also have larger negative effects on employment and larger positive effects on the probability of receiving unemployment benefits.

Together with the heterogeneity by sex, these results suggest that outsourcing affects the distribution of earnings not only between individuals, but also between groups of different migratory background by concentrating its negative effects on groups known to be more vulnerable on the labour market.

4.3 Robustness checks

Including small transfers The individuals interviewed in the course of our fieldwork frequently mentioned small cleaning or collective catering teams comprising as few as three or four workers. This suggests that the transfer definition adopted in our main specification, which requires at least six employees, is restrictive, excluding many smaller teams. That relatively high threshold guards against false positives (cases where, by chance, several workers who leave the same establishment A

in a given year happen to join the same subcontractor B), but it may yield a biased picture of the universe of transfers. As a robustness check, we therefore replicate the analysis on an expanded sample encompassing all transfers of three or more workers, which increases the number of treated workers by 60%, as reported in Table 3.

The results are presented in Appendix G.1: they are very similar to the main specification, with some coefficients gaining in statistical significance owing to the larger sample size.

Without controlling on the dependent variables In our main specification, we control on the pre-treatment (from $t - 1$ to $t - 4$) dependent variables (log earnings, log days, log hours per day, employment dummy). This ensures not only parallel dynamics of the two groups in these three dimensions before treatment, but also similar levels (cf. Appendix D), which makes the two groups more comparable. However, these advantages must be weighed against the risk of the so-called Nickell bias that may arise when controlling on the lagged outcome – see Nickell (1981), and de Chaisemartin and D’Haultfœuille (2025, p. 122) for a recent discussion in the diff-in-diff context.

To ensure that our results are not driven by that bias, we run the analysis without any controls on the dependent variables – keeping only sex interacted with age and age squared, migrant status, seniority in the establishment, establishment and firm size, and 3-digit industry and 4-digit occupation fixed effects.

Results are presented in Appendix G.2. They are broadly consistent with those of the main specification, although pre-trends for certain variables limit the interpretability of the estimated coefficients. The effect on the number of days worked is larger, reaching -10 log points at $t + 4$ and persisting at that level thereafter. This in turn drives a larger effect on total earnings, which reaches -15 log points at $t + 5$; however, this figure should be interpreted with caution given that the parallel trends assumption is less well satisfied here, with a statistically significant positive coefficient at $t - 3$.

Notably, no significant effect on the hourly wage is detected in this specification, despite the absence of any pre-trend. This provides an additional reason to question the effect of outsourcing on the hourly wage in our context, especially when set against the robustness of the effect on total earnings.

Without any controls de Chaisemartin and D’Haultfœuille (2025, p. 118-119) warn researchers against the over-reliance on controls and conditional parallel trends. For that reason, we also run our analysis without any controls.

Results are presented in Appendix G.3. Effects on earnings, days and hours per day are very

similar to those of the previous robustness test (no controls on dependent variables), including a pre-treatment significant coefficient on earnings at $t - 3$ that invites to caution when interpreting the results.

Despite the absence of any pre-trend, no effect on the unemployment benefits dummy is measured, which limits the robustness of that finding.

A massive pre-trend on the hourly wage, as well as substantial pre-treatment estimates on the employment dummy, preclude the interpretability of these effects.

$t - 2$ as reference To account for potential anticipation effects, in Appendix G.4, we test taking $t - 2$ as a reference time instead of $t - 1$, accordingly excluding earnings, days and hours at $t - 1$ from the controls.

Results remain broadly consistent with the main specification, although we measure small but significant coefficients on earnings components at $t - 1$: positive on days, negative on the hourly wage and hours per day, which cancel out into a null effect on earnings. That suggests a modest amount of either anticipation effects or selection into treatment correlated with unobserved characteristics.

5 Interpretation and discussion

5.1 Conditions: the legal framework

Legally, according to articles L1224-1 and -2 of the Labour Code, the transfer of an activity from an employer to another preserves the contract of employment, i.e. all the obligations of the previous employer towards his employees, including wage, hours and employment protection. There are at least two ways for employers to get around this legal constraint. First, many employees do not know their rights or do not feel they are in a position to have them respected – as mentioned in section 2, many of these workers are immigrants and do not have any degree, and union density in these occupations is very low. Second, it appears that a common strategy from employers is to use the so-called *clause de mobilité* (mobility clause) in the employment contract to re-assign some hours of workers from one site to another, potentially far-away. Unlike a direct change in working hours or compensation, a change in the work site within the geographical scope defined by the mobility clause falls under the employer’s prerogatives and an employee’s refusal may be considered a valid cause for termination. As a consequence, employees will often prefer to accept fewer hours rather than to split their hours each week among different sites that are often far away from one another.

For example, Nadia and Youssef, two site-level cleaning supervisors, explain that the mobility clause in the employment contract in these industries gives the employer a lot of leverage. It is confirmed by Sophie, labour law expert in the HR team of a catering company. At some point in her career, Thérèse, a food service worker, went through such a forced move, ordered by the management to cut costs on the previous site.

5.2 Process: cutting hours on site

We do not directly observe the volume of hours that a subcontractor sells to a client on a given site. The effects that we observe on hours and mobility out of the establishment and out of employment could be compatible with the subcontractor firm letting new, perhaps more efficient or cheaper employees move into the site to compensate. We believe that such a substitution, when it exists, is only partial. Indeed, both the existing qualitative evidence and our own interviews with union representatives and with managers strongly point towards a tendency for subcontractor firms to cut the hours budget allocated to a given site in order to decrease the price and thus win the contract and keep it.

In their report for the French Ministry of Labour, Thevenot et al. (2021) conducted multiple interviews in a large cleaning firm and conclude that “the logic of outsourcing is based on a drastic reduction in the volume of work” (p. 87), because of “the pressure exerted through frequent renegotiations and the emphasis on prices is a general trend” (ibid.). Similarly, Bret (2023) reports that in a French university, the outsourcing of cleaners was followed by a sharp reduction in their number from 160 to 130.

Our interviews point in the same direction. Thérèse, catering assistant, experienced it herself when both the cleaning and catering teams of the clinic where she was working were transferred to a subcontractor. Her new employer used the mobility clause mentioned above to move her to another site against her will (the new site was much further away from her home), and dismissed other colleagues for various causes, without replacing any of them. She estimates that the total size of the team went from 10 to 6. Over the course of her subsequent career, she worked exclusively for subcontractors and experienced two other transfers between firms following contract renewals; in one of these, the team size was again reduced from 7 to 5.5 full-time jobs. She also mentions another case that she is intervening in as a union representative, and claims that “all [subcontractor] companies are trying to reduce their teams” and “it’s not just [her own employer], whether it’s [leading firms in the industry] or anyone else, they’re all playing this little game.”

In cleaning, Nadia says: “The client always looks for the cheapest option. The new company that comes in has to make ends meet, so it has to reduce its workforce. So, with [the previous

subcontractor], we had many more employees [on site].”

On the managers’ side, the answers were more guarded, but overall consistent with the narrative sketched above. Henri, HR at cleaning company A in charge of finding new clients in the health sector, argues that when employed in-house, cleaners tend to be “neglected” in these organisations; by contrast, when outsourced, “just taking care of people, supervising them, appointing a manager, creating job descriptions, reorganising things, because people will be more closely monitored, more supervised, more supported, and better trained, quite simply, they will be more efficient.” But he also mentions in passing that sometimes there is “staff who become in surplus in relation to the bid that has been submitted,” who must then be “properly supported,” for example by transferring them to other sites. In the HR team from catering company D, Marc-André mentions the importance of “staffing” (*postage*, i.e. the number of full-time equivalent positions) in the bidding process with the client company, and cites what he considers an extreme example from a competitor who offered to reduce the team on a site to 2.2 FTE. He argues that his own company would never do this, and would “propose something that would be reasonable like 2.6, 2.8 or 3.6 FTE”.

His colleague Sophie, who is the specialist in transfers in the team, also stresses that her company refrains from the “aggressive” bidding practices of some of its competitors, and that “in general, more or less, we will proceed with a similar staffing to the one already in place.” But she did not rule out the existence of such reductions in full-time equivalent staffing, and went on to explain how they would typically be put in place: an adaptation period of 3 to 12 months would be negotiated with the client, and during that period the HR team would try to convince one of the employees on site to move to another site that suits her, or to mutually agree to terminate the contract (*rupture conventionnelle*). The *mobility clause* giving the employer the right to move the employee to another site without her consent, she says, gives them an additional “flexibility”, but “our philosophy favours concerted rather than forced mobility”.

The clearest exception to this trend in our interview data comes from Aïssata, who handles relations with the cleaning service provider at the French branch of a multinational tech company. She explains that they have kept the same cleaning service provider for five years, with little staff turnover among the cleaning employees, and even a recent expansion of the team. She attributes this to the fact that the client company (her employer) is growing rapidly and therefore needs more and more cleaning staff rather than fewer, and to the priority given to the well-being of on-site employees – highly skilled engineers who are attentive to the quality of cleaning and catering services.

5.3 Consequences

Such a decrease, when it happens, must necessarily be a convex combination of the following four factors: 1. an improvement in productivity at a given level of work effort; 2. an intensification of work, i.e. an increase in the effort provided in a given volume of paid hours; 3. unpaid labour hours; 4. a decrease in the quality of the service. In the existing and our qualitative research, we find evidence mostly of (2) and (4).

Intensification Both Nadia (union organiser) and Malika (cleaner) explained that after a contract renewal with the university, the contractor dismissed several workers and this increased the workload of the remaining ones “Interviewer. – When there are these contract renewals entailing a lower budget and fewer staff, what does that mean for the employees? Nadia. – Work overload. Well, it’s a work overload, because the surfaces stay the same, so the work stays the same, the classrooms stay the same. But when staff numbers go down, then in fact the workload increases for those who remain.”

Decrease in quality That a cut in the hours budget on a given site was followed by a decrease in quality, was mentioned in our interviews in both cleaning and catering. “The worker has to clean around 10 offices per day. To do an office properly, that is to say disinfecting the telephone, the chairs, the contact points, the lower finishes, the upper finishes, the bookshelves, it would take him at least 20 minutes per office. But in the schedule we give him, he’s supposed to spend 5 minutes, 6 minutes at most.” (Nadia) “Now they’re asking you to make everything spotless once a week. That’s not possible. The most you can do... well... if you manage to do 60%, that’s already good.” (Youssef) This resonates with the qualitative findings of Zuberi (2013) on the effects of outsourcing hospital cleaning in Canada: according to that study, it led to increasing work intensity and decreasing hygiene, which coincided with an increase in nosocomial infections.

In the catering industry, both Thérèse and Alain mention a decrease in quality following cuts in the total hours on a given site. They both cite a switch from fresh to frozen food as a way to save on hours, and Thérèse also mentions a pressure to spend less time with patients in a hospital restaurant.

Unpaid hours Unpaid hours are illegal, but can result from managerial pressure on the amount of work to do on a given day. The only interviewee who mentioned it was Thérèse, who said that it is frequent at her company, especially following cuts in jobs assigned to a given site and the pressure that ensues for the remaining employees to keep the restaurant running. Other interviewees did not mention it, and Malika explicitly denied that it happened to her, so we are not able to assess

the importance of the phenomenon.

Let us still note that according to several anecdotal accounts, the basic legal principle of compensation proportional to labour time is frequently reneged upon in the cleaning sector. In her account of her participant observation as an outsourced cleaner in Normandy, Aubenas (2010) tells how her team was assigned a quota of camping bungalows to clean in an unrealistic time, and felt compelled to accept because of the fear of losing their hours. One of the main demands of a hotel cleaners' strike in Paris in 2012 was to shift from piece-rate pay to hourly wages (Doumayrou, 2013). See also the testimony of a hotel cleaner in Marseille in 2019 after a transfer of the contract from an employer to another: "In our contract, we are supposed to work 5 hours a day. But in reality, in the morning they assign us 10 or 12 rooms... Sometimes some of them are very dirty, we take more time there. So in general, we get out only at 3pm, or even 4, instead of 2." and according to her, these hours are not paid by the new employer (Hubinet, 2019).

5.4 A Comparison with previous studies

These results bear both a striking similarity to and a notable difference with Goldschmidt and Schmieder (2017)'s study on outsourcing into food, cleaning, logistics and security in Germany between 1975 and 2009, from which we closely replicate the definition of an outsourcing event as well as the propensity score weighting approach to building the control group. They find a penalty in the daily wage of around -5 log points immediately after outsourcing, reaching around -10 log points 10 years after, so an effect similar to ours in the short run and larger in the long run.

However, German matched employer-employee administrative data do not include a number of hours. The authors argue (p. 1186) that the effect they measure is driven by the hourly wage, based on indirect evidence: they do not find any significant difference between the treated and the control group after outsourcing in either days worked or full-time status, and the effect on earnings is not affected by restricting the sample to full-timer workers before treatment.

This divergence from our results can be explained in two potentially complementary ways. On the one hand, part or all the earnings penalty they measure may be driven by a fall in hours not reflected in the actual or at least the reported full-time status⁶. In the second half of their period of study, from the 1990s onwards, this would be consistent with the fact reported by Checchi et al. (2016), that the increase in low pay, part-time jobs played an important role in the overall increase in earnings inequality.

⁶A large fall in hours is likely to be reflected in the share of full-time employees. Indeed, in our own study we measure a large negative effect of around 10 pp on the probability of working 35 hours or more, the legal duration of the workweek. Further inquiry on the reporting practices of German firms would indicate whether they systematically report in the administrative data when a change in a given employee's hours crosses the full-time threshold.

On the other hand, it may be that different mechanisms for the outsourcing earnings penalty arise from different institutional and economic contexts. Indeed, in the period covered, (West-) Germany had no federal minimum wage, with wage floors mainly set at the industry level, and typically higher in the manufacturing sector, which may have been an important source of outsourcing flows, in contrast to our study; importantly, we do not know whether outsourcing transfers were regulated by contract-preserving clauses in the German context as they are in contemporary France (we will provide explanations in section 5 for why these clauses may work better at protecting the hourly wage than the hours worked).

The remarks and open questions above also apply, *mutatis mutandis*, to Gürer and Taymaz (2025)'s event-study of outsourcing in contemporary Turkey between 2012 and 2022. They find a dynamic effect of outsourcing on the daily wage of low-skilled workers that is very similar to Goldschmidt and Schmieder (2017), but cannot decompose between hours and hourly wage as those are not included in their data.

6 Conclusion

Drawing on a reconstructed panel from French matched employer–employee data covering the period 2001–2019, we show that the outsourcing of a team of catering, cleaning, or security workers inflicts on them a substantial and persistent earnings penalty, accounted for almost entirely by the number of hours worked over the year and at most in small part by the hourly wage.

The seminal study by Goldschmidt and Schmieder (2017) on Germany over the period 1975–2009 found an effect of similar magnitude on the daily wage, but no effect on the probability of working full-time, strongly suggesting that the hourly wage was the primary driver. It seems that differences in economic context or institutions, whether across countries or time periods, shape the margins along which subcontracting firms cut costs and, consequently, the nature of the burden borne by affected workers.

We further find a negative effect on the probability of being in employment and a positive effect on the probability of receiving unemployment benefits. Our interviews with various stakeholders in the catering and cleaning industries indicate that reducing hours and headcount on a given site is a common cost-cutting strategy for subcontractors, driven by the intense competition among them that is sustained by the regular renewal of tendering processes. These effects, like that on earnings, are more pronounced among women than among men, and among migrants than among native-born workers. The on-site outsourcing of low-skilled tasks thus contributes not only to widening interpersonal wage inequality, but also to deepening the disadvantage faced by the most vulnerable

groups.

We hope that these findings will stimulate further research aimed at better understanding the mechanisms underlying these effects, which lie at the intersection of labour market dynamics and bargaining relationships between firms. A satisfactory theory should be consistent both with the existence of a substantial outsourcing penalty, and with its heterogeneity across institutional settings and across workers' characteristics such as gender and migration status.

References

- Abadie, A. (2005). Semiparametric difference-in-differences estimators. *The review of economic studies*, 72(1):1–19.
- Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.
- Aeppli, C. (2025). Subcontracting and employment instability.
- Aubenas, F. (2010). *Le Quai de Ouistreham*. Éditions de l'Olivier, Paris.
- Babet, D., Godechot, O., and Palladino, M. G. (2026). In the land of AKM: Explaining the dynamics of wage inequality in France. *Journal of Labor Economics*. forthcoming.
- Beckmannshagen, M. and Schröder, C. (2022). Earnings inequality and working hours mismatch. *Labour Economics*, 76:102184.
- Bergeaud, A., Cahuc, P., Malgouyres, C., Signorelli, S., and Zuber, T. (2024). The wage of temporary agency workers.
- Bergeaud, A., Malgouyres, C., Mazet-Sonilhac, C., and Signorelli, S. (2025). Technological change and domestic outsourcing. *Journal of Labor Economics*.
- Bergman, P., Chetty, R., DeLuca, S., Hendren, N., Katz, L. F., and Palmer, C. (2024). Creating moves to opportunity: Experimental evidence on barriers to neighborhood choice. *American Economic Review*, 114(5):1281–1337.
- Bilal, A. and Lhuillier, H. (2021). Outsourcing, inequality and aggregate output. Technical report, National Bureau of Economic Research.
- Borowczyk-Martins, D. and Lalé, E. (2019). Employment adjustment and part-time work: Lessons from the united states and the united kingdom. *American Economic Journal: Macroeconomics*, 11(1):389–435.

- Bret, H. (2023). L’université: un terrain de lutte pour les agents de nettoyage en sous-traitance? *Mouvements*, 113(1):40–51.
- Brugière, F. (2017). *La sous-traitance en piste: les ouvriers de l’assistance aéroportuaire*. Érés Toulouse.
- Callaway, B. and Sant’Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.
- Card, D., Heining, J., and Kline, P. (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly journal of economics*, 128(3):967–1015.
- Checchi, D., Garca-Peñalosa, C., and Vivian, L. (2022). Hours inequality.
- Checchi, D., García-Peñalosa, C., and Vivian, L. (2016). Are changes in the dispersion of hours worked a cause of increased earnings inequality? *IZA Journal of European Labor Studies*, 5(1):15.
- Cohen, N., Ghio, N., and Gilbert, M. (2025). Wishing to work more? preferences, constraints, and hours worked.
- Colonna, A. and Aldeco Leo, L. R. (2024). Outsourcing, labor regulations and profit-sharing: Evidence from mexico.
- de Chaisemartin, C. and D’Haultfoeuille, X. (2025). *Credible answers to hard questions: differences-in-differences for natural experiments*. SSRN.
- Devetter, F.-X., Valentin, J., Gadrey, J., and Gadrey, N. (2021). *Deux millions de travailleurs et des poussières-L’avenir des emplois du nettoyage dans une société*. Les petits matins.
- Doumayrou, F. (2013). Victoire contre la sous-traitance à l’hôtel campanile. *L’Humanité*. Consulted in Aug. 2025.
- Drenik, A., Jäger, S., Plotkin, P., and Schoefer, B. (2023). Paying outsourced labor: Direct evidence from linked temp agency-worker-client data. *Review of Economics and Statistics*, 105(1):206–216.
- Dube, A., Girardi, D., Jorda, O., and Taylor, A. M. (2023). A local projections approach to difference-in-differences event studies. Technical report, National Bureau of Economic Research Cambridge, Massachusetts.
- Dube, A. and Kaplan, E. (2010). Does outsourcing reduce wages in the low-wage service occupations? evidence from janitors and guards. *ILR Review*, 63(2):287–306.

- Estefan, A., Gerhard, R., Kaboski, J. P., Kondo, I. O., and Qian, W. (2024). Outsourcing policy and worker outcomes: Causal evidence from a mexican ban. Technical report, National Bureau of Economic Research.
- Felix, M. and Wong, M. B. (2024). The reallocation effects of domestic outsourcing. Technical report, Cowles Foundation for Research in Economics, Yale University.
- Godechot, O., Tomaskovic-Devey, D., Boza, I., Henriksen, L. F., Hermansen, A. S., Hou, F., Jung, J., Kodama, N., Křížková, A., Lippényi, Z., et al. (2024). The great separation: Top earner segregation at work in advanced capitalist economies. *American Journal of Sociology*, 130(2):439–495.
- Goldschmidt, D. and Schmieder, J. F. (2017). The rise of domestic outsourcing and the evolution of the german wage structure. *The Quarterly Journal of Economics*, 132(3):1165–1217.
- Gürer, E. and Taymaz, E. (2025). Skill-biased wage effects of domestic outsourcing.
- Hubinet, N. (2019). c’est nous, les femmes de chambre, qui allons gagner: plus de deux mois de grève dans un hôtel quatre étoiles. *Basta!* Consulted in Aug. 2025.
- Labanca, C. and Pozzoli, D. (2022). Constraints on hours within the firm. *Journal of Labor Economics*, 40(2):473–503.
- Lachowska, M., Mas, A., Saggio, R., and Woodbury, S. A. (2023). Work hours mismatch. Technical report, National Bureau of Economic Research.
- Lin, K.-H., Hung, K., Piccone, D., Handwerker, E. W., Dey, M., and Voldoire, T. (2025). The contagion of labor: Linking workplace copresence and occupational mobility patterns. *American Journal of Sociology*, 131(1):200–237.
- Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica: Journal of the econometric society*, pages 1417–1426.
- Patault, B. and Lenoir, C. (2024). Customer capital spillovers: Evidence from sales managers in international markets. *American Economic Journal: Applied Economics*, 16(4):404–443.
- Pearce, L. D. (2012). Mixed methods inquiry in sociology. *American behavioral scientist*, 56(6):829–848.
- Perraudin, C., Petit, H., Thèvenot, N., Tinel, B., and Valentin, J. (2014). Inter-firm dependency and employment inequalities: Theoretical hypotheses and empirical tests on french subcontracting relationships. *Review of Radical Political Economics*, 46(2):199–220.

- Puech, I. (2004). Le temps du remue-ménage. conditions d’emploi et de travail de femmes de chambre. *Sociologie du travail*, 46(2):150–167.
- Roux, B. (2026). In search of working time? hours constraints, firms and mobility. Technical report.
- Schmutz-Bloch, B. and Sidibé, M. (2024). Matching, centrality and the urban network. *Journal of Urban Economics*, 144:103706.
- Schütz, G. (2018). *Jeunes, jolies et sous-traitées : les hôteses d’accueil*. La Dispute.
- Song, J., Price, D. J., Guvenen, F., Bloom, N., and Von Wachter, T. (2019). Firming up inequality. *The Quarterly journal of economics*, 134(1):1–50.
- Thevenot, N., Geymond, M., Perez, C., Perraudin, C., Valentin, J., and Devetter, F.-X. (2021). *Entreprise éclatée et périmètre de la représentation collective*. Technical report, Dares.
- Weil, D. (2014). *The fissured workplace: Why work became so bad for so many and what can be done to improve it*. Harvard University Press.
- Zuberi, D. (2013). *Cleaning up: how hospital outsourcing is hurting workers and endangering patients*. Cornell University Press.

A Data

A.1 Measuring days worked in the year

In the BTS database, each line corresponds to a given individual at a given employer. To approximate the cumulated duration of employment during the year, we use three variables: `DATDEB`, the beginning of the employment spell; `DATFIN`, the end of the employment spell; and `DUREE`, the duration of the employment spell, which can be equal to $\text{DATFIN} - \text{DATDEB} + 1$ if the individual has worked in the establishment continuously, or shorter if the individual had several distinct employment periods in that establishment in that year. Days lost because of a leave of absence (esp. sick leave or maternity leave) are included in `DUREE` if the leave is 4 days or less, not if it is longer.

Given these data, we regroup into “periods” the employment spells that overlap. For each period, we compute the number of days worked as the minimum between the extension of the period (date of end of period - date of beginning of period +1) and the sum of durations of the employment spells of the period. The number of days worked in the year is then computed as the sum of number of days worked of all periods.

A.2 Reconstruction of the panel

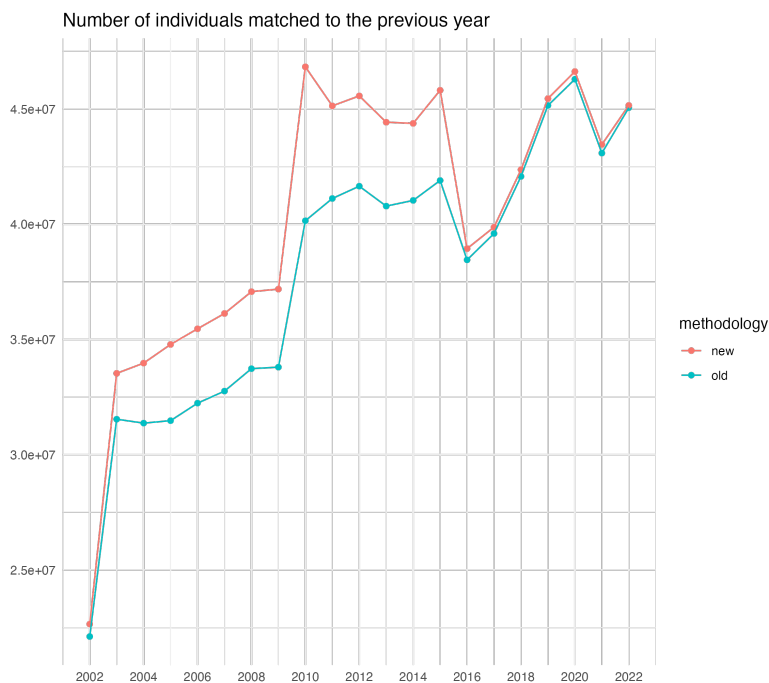


Figure 7: Comparison of the old and the new methodology for chaining BTS year-files

The Insee releases a panel only for a subsample of 8% of individuals. But a given year's files also list, for all individuals appearing in the data in that year, all her employment spells in the previous year. Babet et al. (2026) use this to chain the years between them and, in this way, to reconstruct a quasi-exhaustive panel (see their Appendix C for technical details). That method has been used by a number of authors in recent years (such as Patault and Lenoir, 2024; Schmutz-Bloch and Sidibé, 2024; Lin et al., 2025), and we largely build on it, while improving it on two margins.

First, in the employment spells of year $t - 1$ appearing in the files for year t , the age indicated is the age in year t if the employee has remained in the same establishment, but the age in year $t - 1$ otherwise. The previous algorithm caught any match with any of the two ages, but removed multiples matches. Now we specify correctly the age in each case, which removes some cases of multiple match and thus improves the overall matching rate.

Second, the previous algorithm used only employment spells, while until 2015, the data contain information on unemployment benefits received by persons with or without labour earnings in the

same year. We now take these observations into account, which has two advantages: it allows us to measure the perception of unemployment benefits as an outcome; it allows us to keep track of individuals with a job in year t , without any job but some unemployment benefits in year $t + 1$ (if $t + 1$ is between 2002 and 2015), and a job again in year $t + 2$. With the previous algorithm, the chaining did not chain the jobs in years t and $t + 2$, and would consider them as held by different individuals.

As shown in figure 7, these refinements in the methodology add 2 to 7 millions new matches each year between 2002 and 2015, when the information on unemployment benefits is available, and several hundred thousands for the other years. The matching rate is 97% on average between 2002 and 2019 with the new method and the new domain⁷.

A.3 Sample sizes

The table below reports the size of the treated sample⁸ at time of treatment and 7 years later in the main sample, in several subsamples used in the heterogeneity analysis, and in the larger sample including smaller events (3 or more instead of 6 or more individuals).

Sample	N at t	N at $t + 7$
Main	4,389	1,699
To catering	1,884	990
To cleaning	1,910	511
To security	595	198
To logistics	4,195	1,805
Women	2,899	1,156
Men	1,490	543
Migrants	1,051	313
Natives	3,338	1,386
Balanced panel	2,195	–
Stayers only	3,268	347
Incl. small events	7,074	2,669

Table 3

⁷Note that this figure should not be compared directly to the 98% matching rate claimed for most years by Babet et al. (2026), since the new domain that includes the unemployed is larger.

⁸This is the size of the sample as used in the regression, which requires individuals to be present at $t - 1$ (and, in the “stayers only” specification, to be in the same establishment at the end of $t - 1$ as at the beginning of t), which explains slightly lower figures than in table 1.

B Method

Our method combines the propensity score weighting from Abadie (2005) to define a control group with the local-projection diff-in-diff (LP-DiD) from Dube et al. (2023) to estimate an average treatment effect on the treated. To our knowledge, this combination is novel. We separately propose a package for those interested in using it⁹.

For each outcome y and time horizon h , we run an ordered sequence of three regressions. We denote $T_{t,i} = 0$ or 1 the treatment status (already treated or not treated) of individual i at date t , and $\Delta T_{t,i} = T_{t,i} - T_{t-1,i}$ which is equal to 1 for “newly treated” individuals only, and X_i a vector of individual characteristics that we want to control for. We then define E_h as the set of individuals who, for some t , are observed in the panel at date $t + h$ ¹⁰ and, either are not treated at $t + h$ or were newly treated at t ($\Delta T_{t,i} = 1$ or $T_{t+h,i} = 0$).

First step We run on E_h the propensity score regression:

$$\text{logit}(P(\Delta T_{t,i} = 1 \mid t, X_i)) = \alpha + \delta_t + \gamma X_i \quad (2)$$

This regression yields a predicted probability of treatment $\hat{p}_{i,h}$, from which we derive the inverse probability weights:

$$w_{i,h}^{\text{ipw}} = \begin{cases} 1, & \text{if } \Delta T_{t,i} = 1 \text{ (treated)} \\ \frac{\hat{p}_{i,h}}{1 - \hat{p}_{i,h}}, & \text{if } \Delta T_{t,i} = 0 \text{ and } T_{t+h,i} = 0 \text{ (clean control)} \end{cases} \quad (3)$$

where $\hat{p}_{i,h}$ is the predicted probability of treatment.

Second step Then we run on E_h an auxiliary regression, where each observation is weighted by $w_{i,h}^{\text{ipw}}$:

$$\Delta T_{t,i} = \alpha' + \delta'_t + \varepsilon_{i,t} \quad (4)$$

The residual $\varepsilon_{i,t}$ depends only on the date and the treatment status of the individual. Let us denote $\varepsilon_{t,h}$ its value for units newly treated at date t . If we denote n_t the (propensity-score weighted) number of newly treated individuals at time t , and $N_{t,h}$ the propensity-score weighted number of those who were either newly treated at t or untreated at $t + h$, then Dube et al. (2023,

⁹<https://github.com/oliviergodechot/lpdidcsa>

¹⁰When the outcome considered is the employment dummy or an unemployment benefits dummy, all individuals are considered to be observed in the panel at all dates between 2002 and 2019, either employed or out of employment. For other outcomes, only individuals for which the outcome is directly observed are retained.

Online Appendix, p. 6) note that $\varepsilon_{t,h} = 1 - n_t/N_{t,h}$.

This allows to define new weights as the product of the propensity score weights and cohort weights derived from the ε_t :

$$w_{i,t,h}^* = w_{i,h}^{\text{ipw}} \cdot \frac{\sum_s n_s \varepsilon_{s,h}}{\varepsilon_{t,h}} \quad (5)$$

The first term of the product is used to give more weight to units more similar to the treated group along the chosen dimensions, while the second term is that introduced by Dube et al. (2023) to give the same weight to each cohort and recover in the final stage an “equally-weighted ATT”.

Third step Equipped with these weights, we can run the LP-DiD regression on E_h , where individual i is weighted by $w_{i,t,h}^*$ when considering cohort t at time $t + h$:

$$y_{t+h,i} - y_{t-1,i} = \beta^h \Delta T_{t,i} + \delta_t^h + u_{t,i}^h \quad (6)$$

Then $\widehat{\beta^h}$ is our estimate of the equally-weighted average treatment effect on the treated at horizon h .

C Interviews

We conducted 10 interviews with various stakeholders of subcontractors in cleaning and catering, our two main industries of interest. They lasted between 30 minutes and 1 hour each and were all conducted in France¹¹. Most interviews took place in the Fall 2025, with the only exceptions of Henri (Summer 2023, with a brief follow-up call in March 2026) and Aïssata (March 2026). 4 interviewees were from the cleaning and 5 from the catering industry, and one (Aïssata) works at the French branch of a multinational tech company, where she handles relations with the cleaning service provider ; 5 were managers working in their company’s headquarters, the others’ occupation would have them work on site. 4 of them were elected as union representatives and as such, benefitted from some hours (reaching full-time for 2 of them) off work to the service of the union (*délégation syndicale*). The table below provides a brief summary of the characteristics of each interviewee. All names have been changed.

¹¹Nadia and Youssef, who have been colleagues in the past and are members of the same union, were interviewed jointly.

First name	Origins	Occupation	Industry	Company	Meeting type
Henri	French	Top-level management	Cleaning	A	In person
Youssef	North-African	Site-level supervisor	Cleaning	B	In person
Nadia	North-African	Site-level supervisor	Cleaning	C	In person
Malika	North-African	Cleaner	Cleaning	C	In-person
Marc-André	French	HR manager	Catering	D	In person
Sophie	French	HR manager	Catering	D	Video call
Claire	French	HR manager	Catering	D	In person
Thérèse	West-African	Food service worker	Catering	D	In person
Alain	French Caribbean	Cook	Catering	D	In person
Aïssata	West-African	Cleaning Services Coordinator	Tech	E	Video call

With one exception¹², all non-managers we met are elected union representatives, either at CGT (*Confédération Générale du Travail*, left) or CFDT (*Confédération Française Démocratique du Travail*, centre-left). As such, they benefit from the allocation of part or all of their working time to the service of their union, without a loss in earnings (*délégation syndicale*). This *délégation* is part-time only for Nadia and Alain, but full-time for Youssef and Thérèse.

Companies A, B, C and D are of various sizes, but all have thousands of employees spread across hundreds or more of sites. The clients to which the sites belong are diverse: Nadia and Malika work in a university; Youssef also used to work there, and before also worked as a cleaner in a shopping centre; Alain works in a clinic, as did Thérèse before becoming a full-time union representative.

¹²That of Malika. However, we still met her through the mediation of her supervisor Nadia, a union representative, and they had recently participated in a strike together.

D Averages by group (main specification)




Figure 8: Total yearly earnings in the treated and the control group, in 2015 euros.




Figure 9: Days worked per year in the treated and the control group.




Figure 10: Hours per week in the treated and the control group.




Figure 11: Hourly wage (in 2015 euros) in the treated and the control group.




Figure 12: Share employed in the treated and the control group.




Figure 13: Share receiving unemployment benefits and no earnings in the treated and the control group.




Figure 14: Share receiving unemployment benefits in the treated and the control group.




Figure 15: Share with stable establishment in the treated and the control group.

E Main specification, detailed results




Figure 16: Main specification. Log earnings

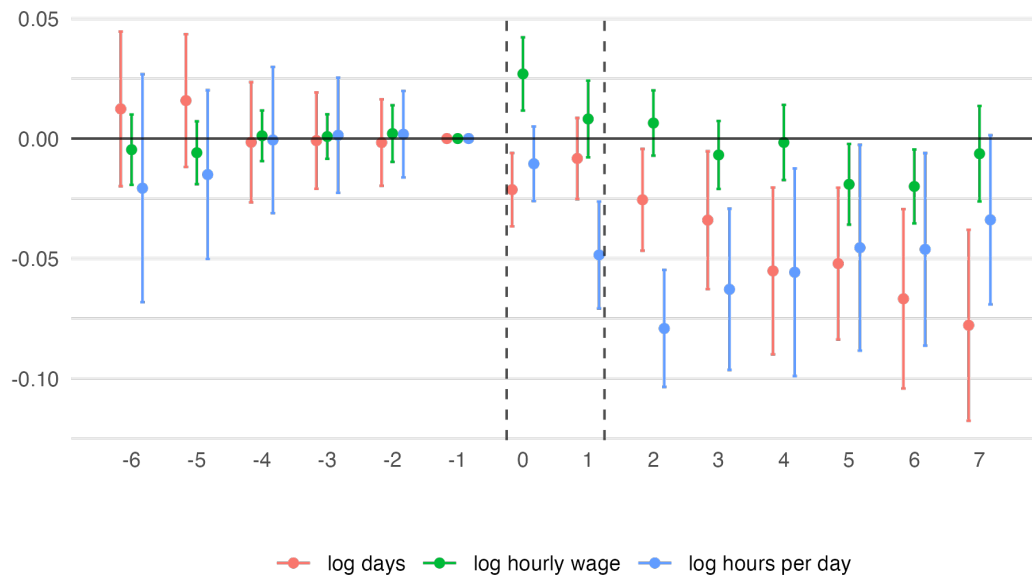


Figure 17: Main specification. Earnings' components

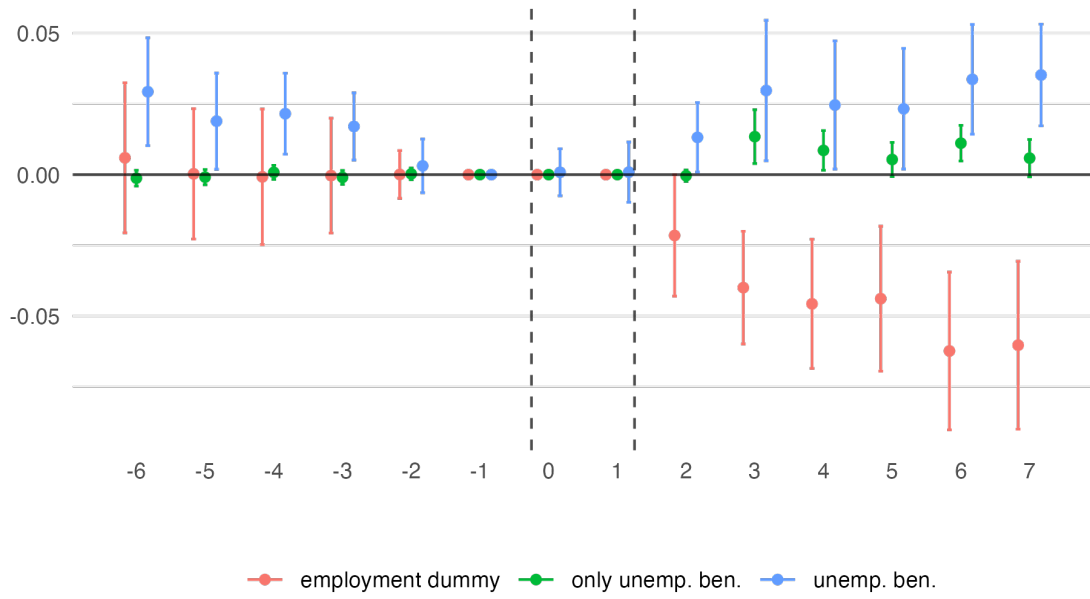


Figure 18: Main specification. Proxies for unemployment




Figure 19: Main specification. Stable establishment dummy

F Heterogeneity analyses

F.1 Balanced panel




Figure 20: Balanced panel. Log earnings




Figure 21: Balanced panel. Earnings' components

F.2 Stayers only




Figure 22: Stayers at the same employer only. Log earnings




Figure 23: Stayers at the same employer only. Earnings' components

F.3 By destination industry

F.3.1 Catering and cleaning




Figure 24: Heterogeneity by industry (catering and cleaning). Log earnings




Figure 25: Heterogeneity by industry (catering and cleaning). Log days




Figure 26: Heterogeneity by industry (catering and cleaning). Log hours per day




Figure 27: Heterogeneity by industry (catering and cleaning). Log hourly wage

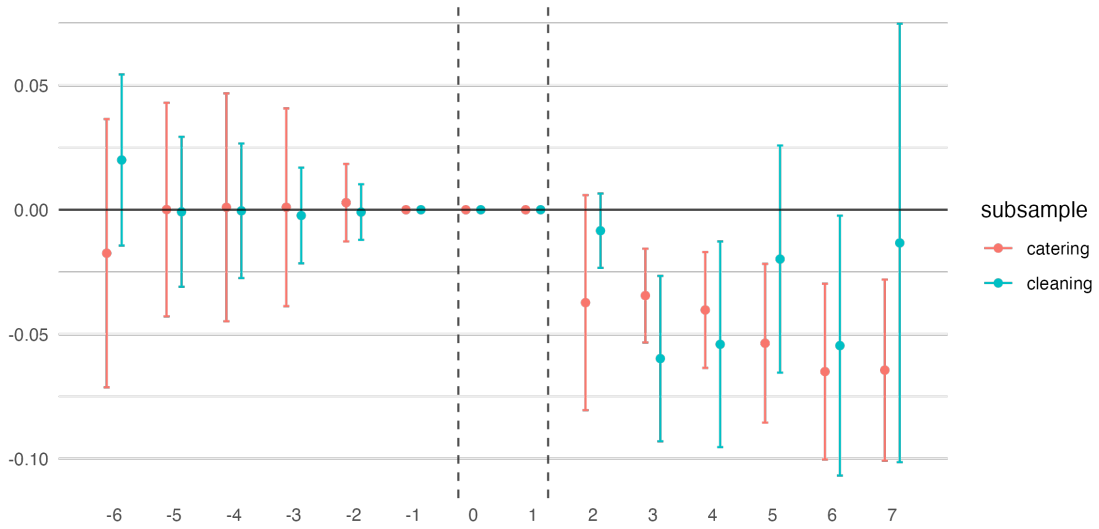


Figure 28: Heterogeneity by industry (catering and cleaning). Employment dummy




Figure 29: Heterogeneity by industry (catering and cleaning). Only unemployment benefits dummy




Figure 30: Heterogeneity by industry (catering and cleaning). Some unemployment benefits dummy




Figure 31: Heterogeneity by industry (catering and cleaning). Stable establishment dummy

F.3.2 Security and logistics




Figure 32: Heterogeneity by industry (security and logistics). Log earnings




Figure 33: Heterogeneity by industry (security and logistics). Log days




Figure 34: Heterogeneity by industry (security and logistics). Log hours per day




Figure 35: Heterogeneity by industry (security and logistics). Log hourly wage




Figure 36: Heterogeneity by industry (security and logistics). Employment dummy




Figure 37: Heterogeneity by industry (security and logistics). Only unemployment benefits dummy




Figure 38: Heterogeneity by industry (security and logistics). Some unemployment benefits dummy




Figure 39: Heterogeneity by industry (security and logistics). Stable establishment dummy

F.4 By sex




Figure 40: Heterogeneity by sex. Log days




Figure 41: Heterogeneity by sex. Log hours per day




Figure 42: Heterogeneity by sex. Log hourly wage




Figure 43: Heterogeneity by sex. Employment dummy




Figure 44: Heterogeneity by sex. Only unemployment benefits dummy




Figure 45: Heterogeneity by sex. Some unemployment benefits dummy




Figure 46: Heterogeneity by sex. Stable establishment dummy

F.5 By migration status




Figure 47: Heterogeneity by migration status. Log days




Figure 48: Heterogeneity by migration status. Log hours per day




Figure 49: Heterogeneity by migration status. Log hourly wage




Figure 50: Heterogeneity by migration status. Employment dummy




Figure 51: Heterogeneity by migration status. Unemployment benefits dummy




Figure 52: Heterogeneity by migration status. Stable establishment dummy

G Robustness checks

G.1 Including small transfers




Figure 53: Including small outsourcing transfers. Log earnings




Figure 54: Including small outsourcing transfers. Earnings' components




Figure 55: Including small outsourcing transfers. Proxies for unemployment




Figure 56: Including small outsourcing transfers. Stable establishment dummy

G.2 Without controls on dependent variables




Figure 57: Without controls on the dependent variables. Log earnings




Figure 58: Without controls on the dependent variables. Earnings' components




Figure 59: Without controls on the dependent variables. Proxies for unemployment




Figure 60: Without controls on the dependent variables. Stable establishment dummy

G.3 Without any controls




Figure 61: Without any controls. Log earnings




Figure 62: Without any controls. Earnings' components




Figure 63: Without any controls. Proxies for unemployment




Figure 64: Without any controls. Stable establishment dummy

G.4 $t - 2$ as reference




Figure 65: $t - 2$ as reference. Log earnings




Figure 66: $t - 2$ as reference. Earnings' components




Figure 67: $t - 2$ as reference. Proxies for unemployment




Figure 68: $t - 2$ as reference. Stable establishment dummy

1 Introduction

There is increasing evidence that inequality between firms affects the labour market and wage distribution. Over the past decades, rising wage inequality between firms has contributed to an increase in total wage variance in Germany (Card et al., 2013) and the United States (Song et al., 2019), while in France, it has offset a compression within firms (Babet et al., 2026).

The same articles note a growing segregation of employees between firms, with lower-paid employees sharing less and less often the same employer with higher-paid employees (see Godechot et al., 2024). These trends can be partly explained by the increasing fragmentation of firms, or ‘fissuring’ (Weil, 2014), which gives rise to asymmetrical commercial relationships between firms. One form of it is temp agency labour, which, as Drenik et al. (2023) and Bergeaud et al. (2024) have shown on Argentinian and French data, partially excludes workers from the rent-sharing at play within companies.

Another important form, which is the subject of this paper, is the on-site outsourcing of low-skilled tasks, such as cleaning, catering and security. We examine that phenomenon in the French context. First, we use the Labour Force Survey (*Enquête Emploi*) to show descriptively that while these industries are small in total employment, they are important for some vulnerable segments of the workforce such as immigrant women without a degree; and that their employees tend to receive low earnings compared to other industries, due to both a low hourly wage and low hours, with an important share of them declaring that they would like to work more hours at the same hourly rate.

We then turn to our main event-study exercise. We use an improved version of Babet et al. (2026)’s procedure to chain administrative matched employer-employee data (*Déclarations Annuelles de Données Sociales*, DADS) into a panel for the period 2001-2019, which allows us to follow Goldschmidt and Schmieder (2017)’s definition of an outsourcing event as a group of workers moving at the same time from a common employer in some other industry to a new one in cleaning, catering or security.

With propensity score reweighting, we define a control group of non-treated or not-yet-treated employees who were similar in terms of observed demographic characteristics, worked in the same origin industry and occupation, and whose earnings and its components were similar in the four preceding years; we then estimate the dynamic treatment effect from this staggered treatment with the Local Projection Diff-in-Diff method from Dube et al. (2023).

We find a penalty of around 10 log points in yearly earnings for treated workers compared to the control group, that persists 7 years after treatment. In contrast to the findings of Goldschmidt and Schmieder (2017), that effect is not mostly driven by the hourly wage. We find an hourly

wage penalty starting two years after treatment, but it is small (1 to 2 log points), explaining only a minor part of the overall earnings effect, and it is significant only for some years and in some specifications.

The effect on earnings is almost entirely explained by a penalty in hours worked per year, both in terms of hours worked per day employed and of number of days worked. After treatment, the average gap between the two groups is of 1 to 1.5 hours per week on the former dimension, and of 3 to 10 days per year on the latter. Heterogeneity analyses reveal that, in the short term, the penalty is driven by workers who voluntarily or involuntarily left the destination employer after the transfer; however, in the medium term, stayers experience a penalty similar to the average. The analyses also show that the penalty is larger for women than for men, and for migrants than for natives.

The penalty in days worked per year can be interpreted as a negative effect on employment. It is complemented by two other metrics: a positive effect of 0.5 to 1 percentage point on the probability of receiving unemployment benefits but no earnings; and a negative effect on the probability of having any employment during the year. Again, these effects are stronger for women than for men, and for migrants than for natives.

The effects on hours and employment are, to the best of our knowledge, a novel finding in the economics literature on the effects of outsourcing. We suggest an interpretation based on 10 interviews conducted with managers, employees and union representatives at subcontractors in cleaning and catering. First, they uncover the role of the so-called “mobility clause”, that allows subcontracting employers to re-assign their employees across large geographical zones, shaping their bargaining power and allowing them to cut hours and jobs. Second, they show the importance of such site-level cuts in hours and jobs in the bidding process in which several subcontractors compete to obtain the contract from a client. Finally, they suggest that the consequences of these cuts are shared between employees (through work intensification) and the principal company or its customers (through a decrease in quality).

Literature

This paper stands at the intersection of two literatures. First, it contributes to the literature on the causes and effects of outsourcing. An early correlational study is that of Dube and Kaplan (2010) who find a wage penalty of 4% or more for outsourced janitors and 8% or more for outsourced guards. The seminal paper for an event study of low skilled service outsourcing is Goldschmidt and Schmieder (2017) on Germany, who find a 10 to 15 % daily wage penalty for workers outsourced into food, cleaning, security and logistics establishments. Applying a similar method on Turkish data,

Gürer and Taymaz (2025) find a stark contrast between low-skilled and high-skilled outsourcing: workers experience a penalty under the former, but a premium under the latter. Colonna and Aldeco Leo (2024) and Estefan et al. (2024) both study outsourcing in Mexico: they show that it was a way to avoid a within-company profit-sharing law, and that an outsourcing ban in 2021 increased compensation for insourced workers. In Brazil, Felix and Wong (2024) study the 1993 lift of an outsourcing ban and find substantial reallocation effects on security guards, where older in-house employees were fired and experienced persistent earnings losses while the task was contracted out to younger employees who benefited.

Several papers have studied outsourcing in France. Matching the REPOSE firm survey with DADS administrative wage data across all industries, Perraudin et al. (2014) show that the subcontractor status of the employer predicts lower wages for observationally similar employees; with the same data sources, Aeppli (2025) shows that it also predicts more employment instability. The causes of outsourcing are still understudied, although Bergeaud et al. (2025) show that when broadband internet reached a new city in France during the period 1999-2007, domestic outsourcing increased in that city, with negative effects on low-skilled and positive effects on high-skilled workers. Bilal and Lhuillier (2021) study service outsourcing in France, both low-skilled (security, cleaning, food, interim, general administrative services, call centres) and high-skilled (accounting, law or consulting services). Running a worker- and firm-fixed effect regression *à la* Abowd et al. (1999) on the (restricted) DADS panel between 1996 and 2007, they find that firm-level wage premia are 14 % lower at contractors compared to non-contractors. They also find that compared to non-contractors, contractors hire more from non-employment and experience a higher separation rate.

There is also a rich sociological and socio-economic literature. Weil (2014) provides a qualitative picture of the effects on workers of various forms of fissuring in the US, including outsourcing but also franchising. In France, outsourcing has been studied in the context of reception hostesses (Schütz, 2018), airport assistance (Brugière, 2017) and cleaning services (Puech, 2004; Devetter et al., 2021; Thevenot et al., 2021). In Canada, Zuberi (2013) studied the effects of the outsourcing of hospital cleaning on workers and on hygiene.

The second strand of literature to which we aim to contribute focuses on part-time work and its determinants. Studies show that part of the increase in German earnings inequality is explained by the increased prevalence of low hours at the bottom of the distribution (Checchi et al., 2016, 2022), while Beckmannshagen and Schröder (2022) use survey questions to show that much of it is involuntary part-time. Lachowska et al. (2023) analyse mobility patterns in the US and infer that many employees, especially at the bottom of the distribution, have a substantial willingness to pay part of their hourly wage for longer hours. Roux (2026) provides similar results on France,

complementing the more descriptive findings of Cohen et al. (2025). Other studies attempt to disentangle the sources of these constraints: while Borowczyk-Martins and Lalé (2019) show on US data that employers reduce workers’ hours during downturns to adjust to falling demand, Labanca and Pozzoli (2022) find evidence in Danish data of both constraints stemming from coordination between employees and from the firm’s “technology”. We contribute to this literature by showing how employers in outsourced service industries reduce workers’ hours to lower their clients’ costs.

Overview

In section 2, we provide descriptive statistics and the institutional context for low-skilled on-site service outsourcing in contemporary France. In section 3, we present our data and estimation strategy. In section 4, we turn to our estimates of the effects on earnings and its components and several proxies for unemployment, with special attention to the heterogeneity of these effects. Finally, in section 5, we interpret these results in the light of the existing qualitative evidence and the interviews we conducted.

2 Context and descriptive statistics

The rise of outsourcing across industries, and in particular in low-skilled services, has been documented in several advanced economies. In the US, between 1984 and 1999, Dube and Kaplan (2010) show an increase in the share of outsourced janitors from 16 % to 22 %, and of outsourced security guards from 40 % to 50 %. In Germany, Goldschmidt and Schmieder (2017) show a dramatic rise between 1975 and 2008, from less than 10% of cleaning and security jobs to almost 30% in security and almost 40% in cleaning.

In France, Bilal and Lhuillier (2021) show, based on the *Enquête Annuelle d’Entreprise* (a large annual firm survey) that outsourcing expenditures as a fraction of aggregate payroll increased from 6% in 1996 to over 10% in 2007, and, based on DADS, that the share of outsourced workers in occupations related to security, cleaning, food, interim, general administrative services and call centers rose by 16 pp, from less than 20 to more than 30% over the same period. Based on the *Enquête Emploi*, Devetter et al. (2021) show that the share of outsourced cleaners doubled from around 25 % in the early 1980s to more than 50 % in 2007, and fluctuating around that level ever since.




Figure 1: Share of the cleaning, catering and security industries among employees. Source: Enquête Emploi, 2013-2019.




Figure 2

As shown by fig. 1, according to the French Labour Force Survey, these three industries together account for just 2 % of the employees, but that weight is much more important for certain

demographic groups, rising to 9 % among women without any degree (despite the overwhelmingly male workforce in security), and even to 16 % for immigrant women without any degree.

As shown in figure 2, according to the 2019 edition of the French Labour Force Survey, the earnings (here, the net earnings from their main job) of employees in the cleaning, catering and security industries are low — at 1390 euros per month, which is two-thirds of the average French employee. That gap is explained by both a lower hourly wage and fewer hours worked per week: 31.4 hours on average in these industries, substantially below the French legal working week of 35 hours, compared to 36.8 hours for the average French employee. The survey allows us to explore this last point further: 21% of respondents employed in these industries declare that the part-time is not voluntarily, but because their employer requires them to; if everyone worked their desired number of hours at their current hourly rate, the average working week in these industries would be 34.4 hours, a 10% increase.

This point can be illustrated by quoting the interview of Malika (see Appendix C), single mother of three, who earns 1000 to 1200 euros by working 22 hours per week, shared among two cleaning employers. A recent increase of about 10% in her hourly wage was “not much”, “not a change”, as contrasted with the months where she gets called for an extra 7-hour shift on sundays – then “I get more hours, I am lucky, I get 1500”.

Of course, a substantial part of these gaps can be explained by the characteristics of these workers, who are among the most vulnerable in the French labour market. 74% of the employees in those industries are not high school graduates (vs. 37% in other industries), 61 % are women (vs 50 %), and 30 % are immigrants (vs. 10%). All these factors contribute to limiting the options available for them on the labour market, as illustrated by our interviews in the HR team of a catering subcontractor. Claire, in charge of recruitment, highlights the ease of filling a food service assistant position (*employées de restauration*, jobs that do not require any formal qualification and not necessarily language skills): they receive “up to 600 or 700 applications” and never less than 10, so that the vacancy is “sometimes filled within a week”. She contrasts it with the difficulty in recruiting for other working-class occupations that require a diploma –, such as chefs, who need a vocational qualification (*CAP, Certificat d’Aptitude Professionnelle*), or, in the industrial company where Claire used to work before, skilled blue-collar workers who needed another, slightly higher qualification (*bac professionnel*). Her colleague Marc-André, in charge of training, stressed the limited French language skills of many food service assistants as restricting their career opportunities, particularly because higher positions in the kitchen require the ability to read instructions in order to apply hygiene guidelines.

A common feature of these three industries is the role played by calls for tenders, where the

principal periodically puts subcontracting companies in competition with one another to obtain the service at the best price. This leads to frequent substitutions of one subcontractor for another, governed by a provision common to the three industry-wide collective agreements¹: when the contract for a given site is awarded to a new provider, the employment contract of the outsourced workers on that site must also be transferred. We will come back to that feature in the interpretation of our quantitative results in section 5.

3 Method

3.1 Data

The original data source is BTS (*Base Tous Salariés* or Database of All Employees, formerly known as DADS, *Déclaration Annuelle de Données Sociales* or Annual Report of Social Data), i.e. wage declaration of all employment spells in France to Social Security. The Insee releases a “panelised” version for 8% of the sample only. But each year’s data also contains information on the previous year, which allowed Babet et al. (2026) to chain them and reconstruct a quasi-exhaustive panel for the period 2002-2019. The matching between annual files is based on establishment, gender, number of hours, start and end dates of the job, place of work and residence, earnings, and age. As explained in more detail in Appendix A.2, we build on their method while improving it on the margin by refining the matching on age and taking unemployment benefits into account. The matching rate on the enlarged domain is 97% on average between 2002 and 2019.

We construct:

- Total earnings in the year as the sum of the earnings of each employment spell (`S_BRUT`) of the individual in the year. These are gross or “posted” earnings, net of employer but not of employee Social Security Contributions. We express them in constant 2015 euros, using the Consumer Price Index from Insee.
- Days worked in the year by aggregating the durations of employment spells, taking overlaps into account (cf. Appendix A.1)
- Daily working hours as the ratio of total hours worked (sum of hours worked in each employment spell `NBHEUR`) to number of days worked. (Note that weekends included within a continuous employment spell are counted as days worked; multiplying by 7 converts this variable into hours per week, as in Figure 9.)

¹Cleaning: IDCC 3043, article 7. Catering: IDCC 1266, Avenant n° 3 du 26 février 1986. Security: IDCC 1351, Accord du 5 mars 2002 relatif à la reprise du personnel.

- Hourly wage as the ratio of total earnings to total hours.
- An employment dummy, which is equal to 1 when the individual has earnings in that year. In other words, an individual is considered as non-employed if she either appears in the panel solely as an unemployment benefit recipient, or does not appear at all.
- A stable establishment dummy. At time $t + h$, an individual is considered to be in the same establishment (as right after treatment) if her main establishment on Dec. 31st is the same as on Dec. 31st $t + 1$. At time $t - h$, an individual is considered to be in the same establishment (as right before treatment) if her main establishment on Jan. 1st is the same as on Jan. 1st t .
- A dummy for receiving unemployment benefits. Note that this information is reported in the data only between 2002 and 2015. Furthermore, unemployment benefits are reported only for individuals who would otherwise be present in the database, i.e. those with some paid employment the same year or the previous year.
- A dummy for receiving *only* unemployment benefits, i.e. receiving unemployment benefits but no earnings in the year.

3.2 Definition of the treatment and the control group

We identify an outsourcing event as a group of at least six persons who are employed on 1 January in the same non-FCSL establishment and on 31 December of the next year in the same catering, cleaning, or security establishment. The outsourced workers need to represent strictly less than half of the original establishment’s workforce, and the original establishment should still exist the next year. That destination establishment should not belong to the same firm or the same business group as the original establishment. We restrict ourselves to workers employed in the private sector the year before treatment².

With this definition, we identify 764 outsourcing events happening between 2005 and 2019 covering 4,558 treated workers³. Most of them were outsourced into catering or cleaning, and only a few hundreds outsourced into security. These events were spread over our whole period of study (2005-2019), although irregularly, with a maximum of 119 events in 2013 and a minimum of 25 events in 2009. Hospitals (by construction, private ones) are the most frequent industry of origin, common to 39 % of treated employees.

²The main establishment in the year before treatment should be either “Organismes privés spécialisés et groupements de droit privé” or “Autres sociétés privées”. This excludes public establishments, “entreprises individuelles” and “particulier employeur”.

³Because, in the main specification, the control group is built to be similar in terms of employment for $t - 1, \dots, t - 4$, we retain only events happening after 2005.

	To catering	To cleaning	To security	From a hospital	2005-2009	2010-2014	2015-2019	Total
# events	242	324	198	169	231	330	203	764
# individuals	1,996	1,956	606	1,782	1,705	1,633	1,220	4,558

Table 1: The treated group and its composition

To define an adequate control group, we first select a random subsample of 0.5% of all individuals present at least once in the panel in the period 2001-2019; after removing all observations of an already treated individuals (clean control condition) and imposing the presence of the individual at the beginning of year t and end of year $t + 1$ for comparability with the treated group, we obtain a control sample of 763 thousand observations. We then use propensity score weighting (also called inverse probability weighting, cf. Abadie (2005) and Appendix B below), i.e. we regress a treatment dummy on a set of predictors (controls), from which we estimate for each individual a predicted probability of being treated, and then re-weight individuals in the control group, assigning a larger weight to those with a larger predicted probability of treatment, i.e. more similar to the treated.

In our main specification, we control for, on the one hand, the range of variables at $t - 1$ – sex interacted with age and age squared, migrant status, establishment and firm size, seniority in establishment (for the main job), number of jobs in the year, 3-digit industry, 4-digit occupation –, and on the other hand, the following variables at each date between $t - 1$ and $t - 4$ included: employment status, log yearly earnings, log days and log hours per day. For reasons of comparability, we impose that members of the control group should work in a private establishment at $t - 1$ and, because that criterion is indirectly used in the definition of the treatment, that they should be employed on 1 January of t and on 31 December of $t + 1$.

1,030 individuals contribute to 10 % of the resulting reweighted control group and 11,343 individuals to 50 % of it⁴. Some characteristics of the treated and the control group are presented in table 2.

3.3 Estimation of the treatment effect

To estimate a treatment effect while avoiding the issues with the naive two-way fixed-effect regression applied in a staggered treatment setting, we combine the propensity-score weighting described

⁴To speak of “the” control group is a slight abuse of language. Here we are referring to the controls at horizon $h = 0$, but as detailed in Appendix B, a control group is defined for each horizon h : when moving from h to $h' > h$, the control group changes both because some individuals who were in the control group at h get treated and thus cannot serve as “clean” controls any more, and because some individuals who were in the treated group at h disappear from the panel at h' , changing the composition of the treated group and thus requiring new weights for the controls.

	Treated	Control
Age	44	44
Female	66 %	67 %
Migrant	24 %	24 %
Est. size	661	680
Firm size	3,087	3,275
Full time	55 %	55 %

Table 2: Average characteristics of the treated and control group the year before treatment

above with the LP-DiD method (Dube et al., 2023), a flexible and computationally light variant of Callaway and Sant’Anna (2021).

The equation estimated for each outcome y (log earnings, log days, log hours per day, log hourly wage, employment dummy, unemployment benefits dummy and stable establishment dummy) and horizon $h \in \{-6, \dots, 7\}$ is as follows:

$$y_{i,t+h} - y_{i,t-1} = \beta_h \Delta T_{i,t} + \delta_t^h + u_{i,h} \quad (1)$$

where y_i is the outcome considered, $\Delta T_{i,t}$ is an indicator for newly treated units at time t , δ_t^h is a year fixed-effect and $u_{i,h}$ is a random error term. All standard errors are clustered at the level of the firm identifier of the main job at $t - 1$.

As required by the method (Dube et al., 2023, p. 6), for β_h to yield an estimator of the equally-weighted average treatment effect on the treated (ATT), equation 1 is estimated each time with “clean controls”, i.e. removing from the sample the “forbidden comparisons” which are at the root of the distortions of the TWFE estimation in a staggered treatment context, and observations are, in addition to the propensity score weights mentioned above, re-weighted by a function of the proportion of treated units in each cohort (see Appendix B).

3.4 Complementary interviews

To complement our quantitative analysis, we conducted nine face-to-face interviews with stakeholders in the cleaning and catering industries, including managers, employees, and union representatives (see Appendix C for details). Initial interviewees were contacted through personal networks, and additional participants were recruited through snowball sampling. While this small convenience sample is not statistically representative, we aimed to include at least one representative from each key functional role (including workers, supervisors, employers, and union delegates) and sector (cleaning and catering). This approach allowed us to cross-reference experiences and perspectives across functional, sectoral, and class divides. The qualitative insights help illuminate

the mechanisms behind our quantitative findings, particularly how cost-cutting strategies shape workers' outcomes. We thus follow both the tradition of mixed-methods research in the social sciences (Pearce, 2012) and recent trends in economics, where qualitative insights help disentangle the mechanisms underlying causal estimates (Bergman et al., 2024).

4 Results

4.1 Average effects

Earnings and its components As shown by figure 3, we find a large and persistent effect on total yearly earnings of close to -10 log points at $t + 2$, reaching -13 log points at $t + 6$. The same figure shows that this earnings penalty is mainly explained by negative effects on the number of hours worked per day and days worked in the year. The former is more important in the short term and the latter in the longer term, although both effects are present and significant over the whole post-treatment window, with only one exception (the effect on hours per day ceases to be significant at $t + 7$). In contrast, the hourly wage explains at most a minor part of the effect on total earnings. Indeed, as shown in figure 17, we measure a significant effect on the hourly wage only for two years, $t + 5$ and $t + 6$, and even then, its magnitude of around -2 log points represents only a minor part of the total effect.

As detailed in section 4.3, that qualitative picture of a substantial and persistent fall in yearly earnings, with at most a small contribution of the hourly wage, is robust to several alternative specifications.

As discussed in section 5.4, that result stands in contrast to the seminal outsourcing event study from Goldschmidt and Schmieder (2017) on German data, which is all the more surprising since we borrowed from them the method for identifying outsourcing transfers in panel earnings data. Indeed, they measure a penalty in the daily wage and, although their data do not allow them to decompose it as we do, they give plausible reasons to interpret it as driven by the hourly wage.

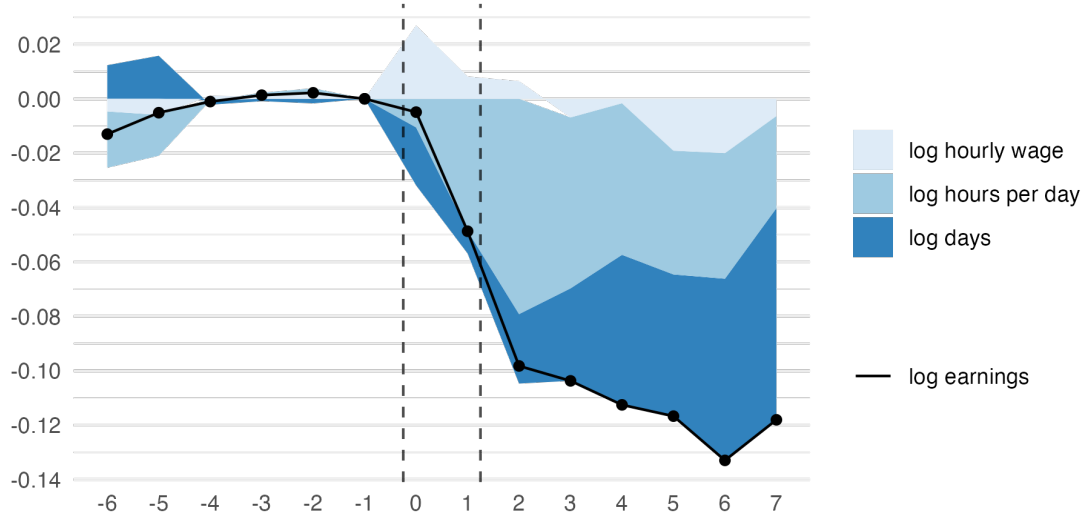


Figure 3: Effect of outsourcing on log earnings and its components. Main specification.




Figure 4: Effect on the probability of employment (sample restricted to those aged 50 or less at time of treatment) and of receiving unemployment benefits (sample restricted to 2002-2015)

Employment effects The penalty in the number of days worked in the year that we just mentioned may reflect both leaves of absence (parental leave and, perhaps more importantly here, sick

leave) as well as additional periods without employment.

We complement it by measuring the effect of treatment on two other outcomes: the probability of having any employment during a calendar year, and the probability of receiving unemployment benefits but no earnings during a year.

These measures have limitations. The main limitation of the employment dummy is related to the very nature of the database used, where an individual appears in year t only if she has a job in year t or had one in year $t - 1$. Because of the way we construct the panel, as detailed in Appendix A, that implies that we lose track of an individual if she is out of employment for a full calendar year and does not receive unemployment benefits during that year, or if she is out of employment for two consecutive calendar years, even if she receives unemployment benefits during these years. Because all units in the sample are employed at t and $t + 1$ by construction, that means that the employment dummy variable is reliable at $t + 2$ but relatively less reliable at longer time horizons, when individuals coming back into employment after one or more years are missed. Note two other limitations in terms of sample: to remove retirement effects, when measuring the effect on employment, we restrict the sample to individuals aged 50 or less at time of treatment; the unemployment benefits are missing from our data after 2015.

With these caveats in mind, the effects reported in Figure 4 are strongly suggestive of negative effects on employment. We find a significant negative and growing effect on the employment dummy, at around -2 pp in the short term and more than -5 pp at $t + 6$ and $t + 7$. We also find a positive effect of around $+1$ pp on the probability of receiving unemployment benefits, significant at $t + 3$, $t + 4$ and $t + 6$. The latter effect is small in the absolute, but high relative to the control group: as shown in figure 14, in some post-treatment years the share of individuals receiving unemployment benefits is actually twice as high in the treated as in the control group⁵.

A further finding deserves mention here: as shown in figure 19, we find a massive negative effect on the probability of keeping the same employer among those who remain in employment, of -20 pp at $t + 3$ and even larger in the longer term. This effect is, however, difficult to interpret, because we cannot isolate the contribution of each of the two underlying mechanisms that it may reflect. One is worker separations, i.e. employees voluntarily or involuntarily leaving their position after being outsourced and subsequently securing employment elsewhere, with the intervening spell of non-employment being consistent with the reduction in annual days worked documented above.

⁵As shown in figure 18, we also compute the effect on the probability of receiving unemployment benefits (including for those with earnings in the same year), but an important pre-trend makes the coefficients uninterpretable. However, when restricting the analysis to either catering or cleaning (fig. 30), the pre-trend disappears and a positive and significant effect remains, between $+2$ and $+5$ pp in catering in years $t + 2$ to $t + 7$, and between $+5$ and $+7$ pp in cleaning in years $t + 4$ to $t + 6$.

The second channel involves employer transitions without job transitions, when the worker’s employment contract is reassigned to an incoming service provider upon contract renewal — a common occurrence in subcontracting markets, as confirmed in our interviews — such that the employee changes employer while remaining in the same role.

4.2 Heterogeneity analysis

Balanced panel Since the share of individuals remaining employed after treatment declines rapidly following treatment (see Figure 12), the effect on earnings at h is estimated on the subsample of treated individuals still employed at that date. This approach makes full use of the available information, but it precludes interpreting the sequence of average effects as the effects experienced by the average treated individual. For this reason, we also estimate, as is common practice, effects within a balanced panel — specifically, a dynamically balanced panel, i.e. retaining only individuals who are continuously observed between $t - 3$ and $t + 5$ (we narrow the time window relative to the main specification to avoid excessive loss of statistical power).

The results are shown in Appendix F.1. We find a significant effect on earnings as early as $t + 2$, but much smaller than in the main specification. It grows over time and exceeds -10 log points by $t + 5$. As in the main specification, this total effect is driven by days worked and hours per day, rather than by the hourly wage, on which we find no significant effect in this specification.

This implies that in the treated group, the most severely affected in the short run tend to subsequently drop out of the panel, with the employment effect thus concentrating — consistent with intuition — among individuals who had already been negatively affected in terms of wages. The remaining individuals, those who remain employed later, experience more moderate short-run effects, but their situation (relative to the control group) then deteriorates over time.

Stayers As noted above, and as can be seen in Figure 15, the share of individuals who remain at the same establishment declines over time, particularly so in the treated group. This means that the average effect on earnings combines the effect on individuals who have stayed employed at the same subcontractor after having been transferred there at the time of treatment, and the effect on individuals who have voluntarily or involuntarily left that job for another.

This second mechanism is of particular interest, as it reveals the practices of subcontractors toward the workforce they take on following a transfer. To isolate it, we run the analysis on a sample that, at each horizon $h \geq 1$, is restricted to individuals who have the same employer at $t + h$ as immediately after treatment (and at $h \leq 0$, to individuals who have the same employer as immediately before treatment; see the description of the stable establishment dummy above in

Section 3.1). Note that this is a conservative definition of stayers: we retain only those who keep the same employer, but not – as they cannot be identified – those who remain at the same site under a different service provider following a contract renewal.

The results are presented in Appendix F.2. In the short run, as in the balanced panel, we find significant negative effects on earnings, but more moderate than in the main specification, indicating that the individuals most affected in the short run are those who immediately left the establishment to which they were transferred. From $t + 4$ onward, however, the effect reaches -10 log points: in the medium run, having been willing or able to remain at the destination establishment of the transfer does not provide protection against a penalty in earnings of roughly the same magnitude as the average. Once again, that penalty is driven by hours worked in the year.

By destination industry We run the analysis separately for each of the three destination industries: catering, cleaning, and security. Results are presented in Appendix F.3.

As mentioned above (Table 1), we identified only 198 outsourcing events towards a security subcontractor, representing 606 individuals; it is therefore probably unsurprising that we do not find any significant effect on this subsample. The comparison between catering and cleaning is more informative. The effects on earnings, on the employment dummy and on the probability of receiving unemployment benefits are broadly similar between catering and cleaning, although the earnings penalty is consistently higher in cleaning. The most striking difference is in the composition of that effect. We find a substantial and significant hourly wage penalty for outsourced cleaners, that explains an important part of the total earnings loss. By contrast, we do not find any such penalty in catering, and we even find a positive, significant effect of $+2$ to $+3$ log points between $t + 2$ and $t + 4$. This means that the null to small negative effect on the hourly wage presented above for the full sample is actually the average between a positive effect in catering and a negative effect in cleaning.

To conclude this paragraph, it is worth mentioning the outsourcing of logistics. Because it is included in the seminal Goldschmidt and Schmieder (2017) paper, we also identified these events and ran a separate analysis on that subsample. Results are in stark contrast to those of our main sample: we do not find any negative effect on earnings, on the employment dummy or on the establishment stability dummy, neither any positive effect on the probability of receiving unemployment benefits. This contrast is maybe not so surprising, given the different characteristics of the industries. Catering, cleaning and security have in common that they involve on-site service provision, which is not always the case in logistics (separate warehouses, transportation); they involve regular renewal of subcontracting contracts, upon which workers are transferred from one subcontractor to another, as stipulated in the collective agreement of each of these three sectors — a mechanism




Figure 5: Heterogeneity by sex. Log earnings

that arguably generates a specific competitive pressure on subcontractors and therefore on workers, which may be less the case in logistics. Still, the result stands in contrast to that of Goldschmidt and Schmiuder (2017), who found a significant and substantial (although milder than in cleaning, see their table II) earnings penalty for outsourced workers in logistics.

By sex As shown in Figure 5, the earnings penalty is consistently larger for women (fluctuating between -10 and -14 log points) than for men (fluctuating between -6 and -9 log points). As shown in Appendix F.4, this is explained by a larger effect both on days and, except at $t + 4$ and $t + 5$, on hours per day.

The contrast on the employment penalty is even more striking. While we find large employment effects on women, ranging from -3 pp at $t + 2$ to -9 pp at $t + 6$, we do not measure any effect on the employment dummy of men. While we accordingly find higher effects on women’s probability of receiving unemployment benefits between $t + 3$ and $t + 4$, this ceases to be true at later dates. Among possible explanations are that these women, although seeking a job, could be less inclined than men to take up their unemployment insurance, or that they may be more likely to leave the labour force.

By migration status As shown in Figure 6, the earnings penalty is consistently larger for migrants than for natives. This is especially true in the medium term, with a massive penalty of




Figure 6: Heterogeneity by migration status. Log earnings

−17 to −19 log points for migrants at $t + 4$, $t + 6$ and $t + 7$, compared to −9 to −12 log points for natives at the same horizons. As shown in Appendix F.5 gap is explained mostly by a more important effect on hours per day and on the hourly wage (along with cleaners, migrants represent another subsample where we measure a significant hourly wage penalty, of −4 log points in the medium term).

Migrants also have larger negative effects on employment and larger positive effects on the probability of receiving unemployment benefits.

Together with the heterogeneity by sex, these results suggest that outsourcing affects the distribution of earnings not only between individuals, but also between groups of different migratory background by concentrating its negative effects on groups known to be more vulnerable on the labour market.

4.3 Robustness checks

Including small transfers The individuals interviewed in the course of our fieldwork frequently mentioned small cleaning or collective catering teams comprising as few as three or four workers. This suggests that the transfer definition adopted in our main specification, which requires at least six employees, is restrictive, excluding many smaller teams. That relatively high threshold guards against false positives (cases where, by chance, several workers who leave the same establishment A

in a given year happen to join the same subcontractor B), but it may yield a biased picture of the universe of transfers. As a robustness check, we therefore replicate the analysis on an expanded sample encompassing all transfers of three or more workers, which increases the number of treated workers by 60%, as reported in Table 3.

The results are presented in Appendix G.1: they are very similar to the main specification, with some coefficients gaining in statistical significance owing to the larger sample size.

Without controlling on the dependent variables In our main specification, we control on the pre-treatment (from $t - 1$ to $t - 4$) dependent variables (log earnings, log days, log hours per day, employment dummy). This ensures not only parallel dynamics of the two groups in these three dimensions before treatment, but also similar levels (cf. Appendix D), which makes the two groups more comparable. However, these advantages must be weighed against the risk of the so-called Nickell bias that may arise when controlling on the lagged outcome – see Nickell (1981), and de Chaisemartin and D’Haultfœuille (2025, p. 122) for a recent discussion in the diff-in-diff context.

To ensure that our results are not driven by that bias, we run the analysis without any controls on the dependent variables – keeping only sex interacted with age and age squared, migrant status, seniority in the establishment, establishment and firm size, and 3-digit industry and 4-digit occupation fixed effects.

Results are presented in Appendix G.2. They are broadly consistent with those of the main specification, although pre-trends for certain variables limit the interpretability of the estimated coefficients. The effect on the number of days worked is larger, reaching -10 log points at $t + 4$ and persisting at that level thereafter. This in turn drives a larger effect on total earnings, which reaches -15 log points at $t + 5$; however, this figure should be interpreted with caution given that the parallel trends assumption is less well satisfied here, with a statistically significant positive coefficient at $t - 3$.

Notably, no significant effect on the hourly wage is detected in this specification, despite the absence of any pre-trend. This provides an additional reason to question the effect of outsourcing on the hourly wage in our context, especially when set against the robustness of the effect on total earnings.

Without any controls de Chaisemartin and D’Haultfœuille (2025, p. 118-119) warn researchers against the over-reliance on controls and conditional parallel trends. For that reason, we also run our analysis without any controls.

Results are presented in Appendix G.3. Effects on earnings, days and hours per day are very

similar to those of the previous robustness test (no controls on dependent variables), including a pre-treatment significant coefficient on earnings at $t - 3$ that invites to caution when interpreting the results.

Despite the absence of any pre-trend, no effect on the unemployment benefits dummy is measured, which limits the robustness of that finding.

A massive pre-trend on the hourly wage, as well as substantial pre-treatment estimates on the employment dummy, preclude the interpretability of these effects.

$t - 2$ as reference To account for potential anticipation effects, in Appendix G.4, we test taking $t - 2$ as a reference time instead of $t - 1$, accordingly excluding earnings, days and hours at $t - 1$ from the controls.

Results remain broadly consistent with the main specification, although we measure small but significant coefficients on earnings components at $t - 1$: positive on days, negative on the hourly wage and hours per day, which cancel out into a null effect on earnings. That suggests a modest amount of either anticipation effects or selection into treatment correlated with unobserved characteristics.

5 Interpretation and discussion

5.1 Conditions: the legal framework

Legally, according to articles L1224-1 and -2 of the Labour Code, the transfer of an activity from an employer to another preserves the contract of employment, i.e. all the obligations of the previous employer towards his employees, including wage, hours and employment protection. There are at least two ways for employers to get around this legal constraint. First, many employees do not know their rights or do not feel they are in a position to have them respected – as mentioned in section 2, many of these workers are immigrants and do not have any degree, and union density in these occupations is very low. Second, it appears that a common strategy from employers is to use the so-called *clause de mobilité* (mobility clause) in the employment contract to re-assign some hours of workers from one site to another, potentially far-away. Unlike a direct change in working hours or compensation, a change in the work site within the geographical scope defined by the mobility clause falls under the employer’s prerogatives and an employee’s refusal may be considered a valid cause for termination. As a consequence, employees will often prefer to accept fewer hours rather than to split their hours each week among different sites that are often far away from one another.

For example, Nadia and Youssef, two site-level cleaning supervisors, explain that the mobility clause in the employment contract in these industries gives the employer a lot of leverage. It is confirmed by Sophie, labour law expert in the HR team of a catering company. At some point in her career, Thérèse, a food service worker, went through such a forced move, ordered by the management to cut costs on the previous site.

5.2 Process: cutting hours on site

We do not directly observe the volume of hours that a subcontractor sells to a client on a given site. The effects that we observe on hours and mobility out of the establishment and out of employment could be compatible with the subcontractor firm letting new, perhaps more efficient or cheaper employees move into the site to compensate. We believe that such a substitution, when it exists, is only partial. Indeed, both the existing qualitative evidence and our own interviews with union representatives and with managers strongly point towards a tendency for subcontractor firms to cut the hours budget allocated to a given site in order to decrease the price and thus win the contract and keep it.

In their report for the French Ministry of Labour, Thevenot et al. (2021) conducted multiple interviews in a large cleaning firm and conclude that “the logic of outsourcing is based on a drastic reduction in the volume of work” (p. 87), because of “the pressure exerted through frequent renegotiations and the emphasis on prices is a general trend” (ibid.). Similarly, Bret (2023) reports that in a French university, the outsourcing of cleaners was followed by a sharp reduction in their number from 160 to 130.

Our interviews point in the same direction. Thérèse, catering assistant, experienced it herself when both the cleaning and catering teams of the clinic where she was working were transferred to a subcontractor. Her new employer used the mobility clause mentioned above to move her to another site against her will (the new site was much further away from her home), and dismissed other colleagues for various causes, without replacing any of them. She estimates that the total size of the team went from 10 to 6. Over the course of her subsequent career, she worked exclusively for subcontractors and experienced two other transfers between firms following contract renewals; in one of these, the team size was again reduced from 7 to 5.5 full-time jobs. She also mentions another case that she is intervening in as a union representative, and claims that “all [subcontractor] companies are trying to reduce their teams” and “it’s not just [her own employer], whether it’s [leading firms in the industry] or anyone else, they’re all playing this little game.”

In cleaning, Nadia says: “The client always looks for the cheapest option. The new company that comes in has to make ends meet, so it has to reduce its workforce. So, with [the previous

subcontractor], we had many more employees [on site].”

On the managers’ side, the answers were more guarded, but overall consistent with the narrative sketched above. Henri, HR at cleaning company A in charge of finding new clients in the health sector, argues that when employed in-house, cleaners tend to be “neglected” in these organisations; by contrast, when outsourced, “just taking care of people, supervising them, appointing a manager, creating job descriptions, reorganising things, because people will be more closely monitored, more supervised, more supported, and better trained, quite simply, they will be more efficient.” But he also mentions in passing that sometimes there is “staff who become in surplus in relation to the bid that has been submitted,” who must then be “properly supported,” for example by transferring them to other sites. In the HR team from catering company D, Marc-André mentions the importance of “staffing” (*postage*, i.e. the number of full-time equivalent positions) in the bidding process with the client company, and cites what he considers an extreme example from a competitor who offered to reduce the team on a site to 2.2 FTE. He argues that his own company would never do this, and would “propose something that would be reasonable like 2.6, 2.8 or 3.6 FTE”.

His colleague Sophie, who is the specialist in transfers in the team, also stresses that her company refrains from the “aggressive” bidding practices of some of its competitors, and that “in general, more or less, we will proceed with a similar staffing to the one already in place.” But she did not rule out the existence of such reductions in full-time equivalent staffing, and went on to explain how they would typically be put in place: an adaptation period of 3 to 12 months would be negotiated with the client, and during that period the HR team would try to convince one of the employees on site to move to another site that suits her, or to mutually agree to terminate the contract (*rupture conventionnelle*). The *mobility clause* giving the employer the right to move the employee to another site without her consent, she says, gives them an additional “flexibility”, but “our philosophy favours concerted rather than forced mobility”.

The clearest exception to this trend in our interview data comes from Aïssata, who handles relations with the cleaning service provider at the French branch of a multinational tech company. She explains that they have kept the same cleaning service provider for five years, with little staff turnover among the cleaning employees, and even a recent expansion of the team. She attributes this to the fact that the client company (her employer) is growing rapidly and therefore needs more and more cleaning staff rather than fewer, and to the priority given to the well-being of on-site employees – highly skilled engineers who are attentive to the quality of cleaning and catering services.

5.3 Consequences

Such a decrease, when it happens, must necessarily be a convex combination of the following four factors: 1. an improvement in productivity at a given level of work effort; 2. an intensification of work, i.e. an increase in the effort provided in a given volume of paid hours; 3. unpaid labour hours; 4. a decrease in the quality of the service. In the existing and our qualitative research, we find evidence mostly of (2) and (4).

Intensification Both Nadia (union organiser) and Malika (cleaner) explained that after a contract renewal with the university, the contractor dismissed several workers and this increased the workload of the remaining ones “Interviewer. – When there are these contract renewals entailing a lower budget and fewer staff, what does that mean for the employees? Nadia. – Work overload. Well, it’s a work overload, because the surfaces stay the same, so the work stays the same, the classrooms stay the same. But when staff numbers go down, then in fact the workload increases for those who remain.”

Decrease in quality That a cut in the hours budget on a given site was followed by a decrease in quality, was mentioned in our interviews in both cleaning and catering. “The worker has to clean around 10 offices per day. To do an office properly, that is to say disinfecting the telephone, the chairs, the contact points, the lower finishes, the upper finishes, the bookshelves, it would take him at least 20 minutes per office. But in the schedule we give him, he’s supposed to spend 5 minutes, 6 minutes at most.” (Nadia) “Now they’re asking you to make everything spotless once a week. That’s not possible. The most you can do... well... if you manage to do 60%, that’s already good.” (Youssef) This resonates with the qualitative findings of Zuberi (2013) on the effects of outsourcing hospital cleaning in Canada: according to that study, it led to increasing work intensity and decreasing hygiene, which coincided with an increase in nosocomial infections.

In the catering industry, both Thérèse and Alain mention a decrease in quality following cuts in the total hours on a given site. They both cite a switch from fresh to frozen food as a way to save on hours, and Thérèse also mentions a pressure to spend less time with patients in a hospital restaurant.

Unpaid hours Unpaid hours are illegal, but can result from managerial pressure on the amount of work to do on a given day. The only interviewee who mentioned it was Thérèse, who said that it is frequent at her company, especially following cuts in jobs assigned to a given site and the pressure that ensues for the remaining employees to keep the restaurant running. Other interviewees did not mention it, and Malika explicitly denied that it happened to her, so we are not able to assess

the importance of the phenomenon.

Let us still note that according to several anecdotal accounts, the basic legal principle of compensation proportional to labour time is frequently reneged upon in the cleaning sector. In her account of her participant observation as an outsourced cleaner in Normandy, Aubenas (2010) tells how her team was assigned a quota of camping bungalows to clean in an unrealistic time, and felt compelled to accept because of the fear of losing their hours. One of the main demands of a hotel cleaners' strike in Paris in 2012 was to shift from piece-rate pay to hourly wages (Doumayrou, 2013). See also the testimony of a hotel cleaner in Marseille in 2019 after a transfer of the contract from an employer to another: "In our contract, we are supposed to work 5 hours a day. But in reality, in the morning they assign us 10 or 12 rooms... Sometimes some of them are very dirty, we take more time there. So in general, we get out only at 3pm, or even 4, instead of 2." and according to her, these hours are not paid by the new employer (Hubinet, 2019).

5.4 A Comparison with previous studies

These results bear both a striking similarity to and a notable difference with Goldschmidt and Schmieder (2017)'s study on outsourcing into food, cleaning, logistics and security in Germany between 1975 and 2009, from which we closely replicate the definition of an outsourcing event as well as the propensity score weighting approach to building the control group. They find a penalty in the daily wage of around -5 log points immediately after outsourcing, reaching around -10 log points 10 years after, so an effect similar to ours in the short run and larger in the long run.

However, German matched employer-employee administrative data do not include a number of hours. The authors argue (p. 1186) that the effect they measure is driven by the hourly wage, based on indirect evidence: they do not find any significant difference between the treated and the control group after outsourcing in either days worked or full-time status, and the effect on earnings is not affected by restricting the sample to full-timer workers before treatment.

This divergence from our results can be explained in two potentially complementary ways. On the one hand, part or all the earnings penalty they measure may be driven by a fall in hours not reflected in the actual or at least the reported full-time status⁶. In the second half of their period of study, from the 1990s onwards, this would be consistent with the fact reported by Checchi et al. (2016), that the increase in low pay, part-time jobs played an important role in the overall increase in earnings inequality.

⁶A large fall in hours is likely to be reflected in the share of full-time employees. Indeed, in our own study we measure a large negative effect of around 10 pp on the probability of working 35 hours or more, the legal duration of the workweek. Further inquiry on the reporting practices of German firms would indicate whether they systematically report in the administrative data when a change in a given employee's hours crosses the full-time threshold.

On the other hand, it may be that different mechanisms for the outsourcing earnings penalty arise from different institutional and economic contexts. Indeed, in the period covered, (West-) Germany had no federal minimum wage, with wage floors mainly set at the industry level, and typically higher in the manufacturing sector, which may have been an important source of outsourcing flows, in contrast to our study; importantly, we do not know whether outsourcing transfers were regulated by contract-preserving clauses in the German context as they are in contemporary France (we will provide explanations in section 5 for why these clauses may work better at protecting the hourly wage than the hours worked).

The remarks and open questions above also apply, *mutatis mutandis*, to Gürer and Taymaz (2025)'s event-study of outsourcing in contemporary Turkey between 2012 and 2022. They find a dynamic effect of outsourcing on the daily wage of low-skilled workers that is very similar to Goldschmidt and Schmieder (2017), but cannot decompose between hours and hourly wage as those are not included in their data.

6 Conclusion

Drawing on a reconstructed panel from French matched employer–employee data covering the period 2001–2019, we show that the outsourcing of a team of catering, cleaning, or security workers inflicts on them a substantial and persistent earnings penalty, accounted for almost entirely by the number of hours worked over the year and at most in small part by the hourly wage.

The seminal study by Goldschmidt and Schmieder (2017) on Germany over the period 1975–2009 found an effect of similar magnitude on the daily wage, but no effect on the probability of working full-time, strongly suggesting that the hourly wage was the primary driver. It seems that differences in economic context or institutions, whether across countries or time periods, shape the margins along which subcontracting firms cut costs and, consequently, the nature of the burden borne by affected workers.

We further find a negative effect on the probability of being in employment and a positive effect on the probability of receiving unemployment benefits. Our interviews with various stakeholders in the catering and cleaning industries indicate that reducing hours and headcount on a given site is a common cost-cutting strategy for subcontractors, driven by the intense competition among them that is sustained by the regular renewal of tendering processes. These effects, like that on earnings, are more pronounced among women than among men, and among migrants than among native-born workers. The on-site outsourcing of low-skilled tasks thus contributes not only to widening interpersonal wage inequality, but also to deepening the disadvantage faced by the most vulnerable

groups.

We hope that these findings will stimulate further research aimed at better understanding the mechanisms underlying these effects, which lie at the intersection of labour market dynamics and bargaining relationships between firms. A satisfactory theory should be consistent both with the existence of a substantial outsourcing penalty, and with its heterogeneity across institutional settings and across workers' characteristics such as gender and migration status.

References

- Abadie, A. (2005). Semiparametric difference-in-differences estimators. *The review of economic studies*, 72(1):1–19.
- Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.
- Aeppli, C. (2025). Subcontracting and employment instability.
- Aubenas, F. (2010). *Le Quai de Ouistreham*. Éditions de l'Olivier, Paris.
- Babet, D., Godechot, O., and Palladino, M. G. (2026). In the land of AKM: Explaining the dynamics of wage inequality in France. *Journal of Labor Economics*. forthcoming.
- Beckmannshagen, M. and Schröder, C. (2022). Earnings inequality and working hours mismatch. *Labour Economics*, 76:102184.
- Bergeaud, A., Cahuc, P., Malgouyres, C., Signorelli, S., and Zuber, T. (2024). The wage of temporary agency workers.
- Bergeaud, A., Malgouyres, C., Mazet-Sonilhac, C., and Signorelli, S. (2025). Technological change and domestic outsourcing. *Journal of Labor Economics*.
- Bergman, P., Chetty, R., DeLuca, S., Hendren, N., Katz, L. F., and Palmer, C. (2024). Creating moves to opportunity: Experimental evidence on barriers to neighborhood choice. *American Economic Review*, 114(5):1281–1337.
- Bilal, A. and Lhuillier, H. (2021). Outsourcing, inequality and aggregate output. Technical report, National Bureau of Economic Research.
- Borowczyk-Martins, D. and Lalé, E. (2019). Employment adjustment and part-time work: Lessons from the united states and the united kingdom. *American Economic Journal: Macroeconomics*, 11(1):389–435.

- Bret, H. (2023). L’université: un terrain de lutte pour les agents de nettoyage en sous-traitance? *Mouvements*, 113(1):40–51.
- Brugière, F. (2017). *La sous-traitance en piste: les ouvriers de l’assistance aéroportuaire*. Érés Toulouse.
- Callaway, B. and Sant’Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.
- Card, D., Heining, J., and Kline, P. (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly journal of economics*, 128(3):967–1015.
- Checchi, D., Garca-Peñalosa, C., and Vivian, L. (2022). Hours inequality.
- Checchi, D., García-Peñalosa, C., and Vivian, L. (2016). Are changes in the dispersion of hours worked a cause of increased earnings inequality? *IZA Journal of European Labor Studies*, 5(1):15.
- Cohen, N., Ghio, N., and Gilbert, M. (2025). Wishing to work more? preferences, constraints, and hours worked.
- Colonna, A. and Aldeco Leo, L. R. (2024). Outsourcing, labor regulations and profit-sharing: Evidence from mexico.
- de Chaisemartin, C. and D’Haultfoeuille, X. (2025). *Credible answers to hard questions: differences-in-differences for natural experiments*. SSRN.
- Devetter, F.-X., Valentin, J., Gadrey, J., and Gadrey, N. (2021). *Deux millions de travailleurs et des poussières-L’avenir des emplois du nettoyage dans une société*. Les petits matins.
- Doumayrou, F. (2013). Victoire contre la sous-traitance à l’hôtel campanile. *L’Humanité*. Consulted in Aug. 2025.
- Drenik, A., Jäger, S., Plotkin, P., and Schoefer, B. (2023). Paying outsourced labor: Direct evidence from linked temp agency-worker-client data. *Review of Economics and Statistics*, 105(1):206–216.
- Dube, A., Girardi, D., Jorda, O., and Taylor, A. M. (2023). A local projections approach to difference-in-differences event studies. Technical report, National Bureau of Economic Research Cambridge, Massachusetts.
- Dube, A. and Kaplan, E. (2010). Does outsourcing reduce wages in the low-wage service occupations? evidence from janitors and guards. *ILR Review*, 63(2):287–306.

- Estefan, A., Gerhard, R., Kaboski, J. P., Kondo, I. O., and Qian, W. (2024). Outsourcing policy and worker outcomes: Causal evidence from a mexican ban. Technical report, National Bureau of Economic Research.
- Felix, M. and Wong, M. B. (2024). The reallocation effects of domestic outsourcing. Technical report, Cowles Foundation for Research in Economics, Yale University.
- Godechot, O., Tomaskovic-Devey, D., Boza, I., Henriksen, L. F., Hermansen, A. S., Hou, F., Jung, J., Kodama, N., Křížková, A., Lippényi, Z., et al. (2024). The great separation: Top earner segregation at work in advanced capitalist economies. *American Journal of Sociology*, 130(2):439–495.
- Goldschmidt, D. and Schmieder, J. F. (2017). The rise of domestic outsourcing and the evolution of the german wage structure. *The Quarterly Journal of Economics*, 132(3):1165–1217.
- Gürer, E. and Taymaz, E. (2025). Skill-biased wage effects of domestic outsourcing.
- Hubinet, N. (2019). c’est nous, les femmes de chambre, qui allons gagner: plus de deux mois de grève dans un hôtel quatre étoiles. *Basta!* Consulted in Aug. 2025.
- Labanca, C. and Pozzoli, D. (2022). Constraints on hours within the firm. *Journal of Labor Economics*, 40(2):473–503.
- Lachowska, M., Mas, A., Saggio, R., and Woodbury, S. A. (2023). Work hours mismatch. Technical report, National Bureau of Economic Research.
- Lin, K.-H., Hung, K., Piccone, D., Handwerker, E. W., Dey, M., and Voldoire, T. (2025). The contagion of labor: Linking workplace copresence and occupational mobility patterns. *American Journal of Sociology*, 131(1):200–237.
- Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica: Journal of the econometric society*, pages 1417–1426.
- Patault, B. and Lenoir, C. (2024). Customer capital spillovers: Evidence from sales managers in international markets. *American Economic Journal: Applied Economics*, 16(4):404–443.
- Pearce, L. D. (2012). Mixed methods inquiry in sociology. *American behavioral scientist*, 56(6):829–848.
- Perraudin, C., Petit, H., Thèvenot, N., Tinel, B., and Valentin, J. (2014). Inter-firm dependency and employment inequalities: Theoretical hypotheses and empirical tests on french subcontracting relationships. *Review of Radical Political Economics*, 46(2):199–220.

- Puech, I. (2004). Le temps du remue-ménage. conditions d’emploi et de travail de femmes de chambre. *Sociologie du travail*, 46(2):150–167.
- Roux, B. (2026). In search of working time? hours constraints, firms and mobility. Technical report.
- Schmutz-Bloch, B. and Sidibé, M. (2024). Matching, centrality and the urban network. *Journal of Urban Economics*, 144:103706.
- Schütz, G. (2018). *Jeunes, jolies et sous-traitées : les hôteses d’accueil*. La Dispute.
- Song, J., Price, D. J., Guvenen, F., Bloom, N., and Von Wachter, T. (2019). Firming up inequality. *The Quarterly journal of economics*, 134(1):1–50.
- Thevenot, N., Geymond, M., Perez, C., Perraudin, C., Valentin, J., and Devetter, F.-X. (2021). *Entreprise éclatée et périmètre de la représentation collective*. Technical report, Dares.
- Weil, D. (2014). *The fissured workplace: Why work became so bad for so many and what can be done to improve it*. Harvard University Press.
- Zuberi, D. (2013). *Cleaning up: how hospital outsourcing is hurting workers and endangering patients*. Cornell University Press.

A Data

A.1 Measuring days worked in the year

In the BTS database, each line corresponds to a given individual at a given employer. To approximate the cumulated duration of employment during the year, we use three variables: `DATDEB`, the beginning of the employment spell; `DATFIN`, the end of the employment spell; and `DUREE`, the duration of the employment spell, which can be equal to $\text{DATFIN} - \text{DATDEB} + 1$ if the individual has worked in the establishment continuously, or shorter if the individual had several distinct employment periods in that establishment in that year. Days lost because of a leave of absence (esp. sick leave or maternity leave) are included in `DUREE` if the leave is 4 days or less, not if it is longer.

Given these data, we regroup into “periods” the employment spells that overlap. For each period, we compute the number of days worked as the minimum between the extension of the period (date of end of period - date of beginning of period +1) and the sum of durations of the employment spells of the period. The number of days worked in the year is then computed as the sum of number of days worked of all periods.

A.2 Reconstruction of the panel




Figure 7: Comparison of the old and the new methodology for chaining BTS year-files

The Insee releases a panel only for a subsample of 8% of individuals. But a given year's files also list, for all individuals appearing in the data in that year, all her employment spells in the previous year. Babet et al. (2026) use this to chain the years between them and, in this way, to reconstruct a quasi-exhaustive panel (see their Appendix C for technical details). That method has been used by a number of authors in recent years (such as Patault and Lenoir, 2024; Schmutz-Bloch and Sidibé, 2024; Lin et al., 2025), and we largely build on it, while improving it on two margins.

First, in the employment spells of year $t - 1$ appearing in the files for year t , the age indicated is the age in year t if the employee has remained in the same establishment, but the age in year $t - 1$ otherwise. The previous algorithm caught any match with any of the two ages, but removed multiples matches. Now we specify correctly the age in each case, which removes some cases of multiple match and thus improves the overall matching rate.

Second, the previous algorithm used only employment spells, while until 2015, the data contain information on unemployment benefits received by persons with or without labour earnings in the

same year. We now take these observations into account, which has two advantages: it allows us to measure the perception of unemployment benefits as an outcome; it allows us to keep track of individuals with a job in year t , without any job but some unemployment benefits in year $t + 1$ (if $t + 1$ is between 2002 and 2015), and a job again in year $t + 2$. With the previous algorithm, the chaining did not chain the jobs in years t and $t + 2$, and would consider them as held by different individuals.

As shown in figure 7, these refinements in the methodology add 2 to 7 millions new matches each year between 2002 and 2015, when the information on unemployment benefits is available, and several hundred thousands for the other years. The matching rate is 97% on average between 2002 and 2019 with the new method and the new domain⁷.

A.3 Sample sizes

The table below reports the size of the treated sample⁸ at time of treatment and 7 years later in the main sample, in several subsamples used in the heterogeneity analysis, and in the larger sample including smaller events (3 or more instead of 6 or more individuals).

Sample	N at t	N at $t + 7$
Main	4,389	1,699
To catering	1,884	990
To cleaning	1,910	511
To security	595	198
To logistics	4,195	1,805
Women	2,899	1,156
Men	1,490	543
Migrants	1,051	313
Natives	3,338	1,386
Balanced panel	2,195	–
Stayers only	3,268	347
Incl. small events	7,074	2,669

Table 3

⁷Note that this figure should not be compared directly to the 98% matching rate claimed for most years by Babet et al. (2026), since the new domain that includes the unemployed is larger.

⁸This is the size of the sample as used in the regression, which requires individuals to be present at $t - 1$ (and, in the “stayers only” specification, to be in the same establishment at the end of $t - 1$ as at the beginning of t), which explains slightly lower figures than in table 1.

B Method

Our method combines the propensity score weighting from Abadie (2005) to define a control group with the local-projection diff-in-diff (LP-DiD) from Dube et al. (2023) to estimate an average treatment effect on the treated. To our knowledge, this combination is novel. We separately propose a package for those interested in using it⁹.

For each outcome y and time horizon h , we run an ordered sequence of three regressions. We denote $T_{t,i} = 0$ or 1 the treatment status (already treated or not treated) of individual i at date t , and $\Delta T_{t,i} = T_{t,i} - T_{t-1,i}$ which is equal to 1 for “newly treated” individuals only, and X_i a vector of individual characteristics that we want to control for. We then define E_h as the set of individuals who, for some t , are observed in the panel at date $t + h$ ¹⁰ and, either are not treated at $t + h$ or were newly treated at t ($\Delta T_{t,i} = 1$ or $T_{t+h,i} = 0$).

First step We run on E_h the propensity score regression:

$$\text{logit}(P(\Delta T_{t,i} = 1 \mid t, X_i)) = \alpha + \delta_t + \gamma X_i \quad (2)$$

This regression yields a predicted probability of treatment $\hat{p}_{i,h}$, from which we derive the inverse probability weights:

$$w_{i,h}^{\text{ipw}} = \begin{cases} 1, & \text{if } \Delta T_{t,i} = 1 \text{ (treated)} \\ \frac{\hat{p}_{i,h}}{1 - \hat{p}_{i,h}}, & \text{if } \Delta T_{t,i} = 0 \text{ and } T_{t+h,i} = 0 \text{ (clean control)} \end{cases} \quad (3)$$

where $\hat{p}_{i,h}$ is the predicted probability of treatment.

Second step Then we run on E_h an auxiliary regression, where each observation is weighted by $w_{i,h}^{\text{ipw}}$:

$$\Delta T_{t,i} = \alpha' + \delta'_t + \varepsilon_{i,t} \quad (4)$$

The residual $\varepsilon_{i,t}$ depends only on the date and the treatment status of the individual. Let us denote $\varepsilon_{t,h}$ its value for units newly treated at date t . If we denote n_t the (propensity-score weighted) number of newly treated individuals at time t , and $N_{t,h}$ the propensity-score weighted number of those who were either newly treated at t or untreated at $t + h$, then Dube et al. (2023,

⁹<https://github.com/oliviergodechot/lpdidcsa>

¹⁰When the outcome considered is the employment dummy or an unemployment benefits dummy, all individuals are considered to be observed in the panel at all dates between 2002 and 2019, either employed or out of employment. For other outcomes, only individuals for which the outcome is directly observed are retained.

Online Appendix, p. 6) note that $\varepsilon_{t,h} = 1 - n_t/N_{t,h}$.

This allows to define new weights as the product of the propensity score weights and cohort weights derived from the ε_t :

$$w_{i,t,h}^* = w_{i,h}^{\text{ipw}} \cdot \frac{\sum_s n_s \varepsilon_{s,h}}{\varepsilon_{t,h}} \quad (5)$$

The first term of the product is used to give more weight to units more similar to the treated group along the chosen dimensions, while the second term is that introduced by Dube et al. (2023) to give the same weight to each cohort and recover in the final stage an “equally-weighted ATT”.

Third step Equipped with these weights, we can run the LP-DiD regression on E_h , where individual i is weighted by $w_{i,t,h}^*$ when considering cohort t at time $t + h$:

$$y_{t+h,i} - y_{t-1,i} = \beta^h \Delta T_{t,i} + \delta_t^h + u_{t,i}^h \quad (6)$$

Then $\widehat{\beta^h}$ is our estimate of the equally-weighted average treatment effect on the treated at horizon h .

C Interviews

We conducted 10 interviews with various stakeholders of subcontractors in cleaning and catering, our two main industries of interest. They lasted between 30 minutes and 1 hour each and were all conducted in France¹¹. Most interviews took place in the Fall 2025, with the only exceptions of Henri (Summer 2023, with a brief follow-up call in March 2026) and Aïssata (March 2026). 4 interviewees were from the cleaning and 5 from the catering industry, and one (Aïssata) works at the French branch of a multinational tech company, where she handles relations with the cleaning service provider ; 5 were managers working in their company’s headquarters, the others’ occupation would have them work on site. 4 of them were elected as union representatives and as such, benefitted from some hours (reaching full-time for 2 of them) off work to the service of the union (*délégation syndicale*). The table below provides a brief summary of the characteristics of each interviewee. All names have been changed.

¹¹Nadia and Youssef, who have been colleagues in the past and are members of the same union, were interviewed jointly.

First name	Origins	Occupation	Industry	Company	Meeting type
Henri	French	Top-level management	Cleaning	A	In person
Youssef	North-African	Site-level supervisor	Cleaning	B	In person
Nadia	North-African	Site-level supervisor	Cleaning	C	In person
Malika	North-African	Cleaner	Cleaning	C	In-person
Marc-André	French	HR manager	Catering	D	In person
Sophie	French	HR manager	Catering	D	Video call
Claire	French	HR manager	Catering	D	In person
Thérèse	West-African	Food service worker	Catering	D	In person
Alain	French Caribbean	Cook	Catering	D	In person
Aïssata	West-African	Cleaning Services Coordinator	Tech	E	Video call

With one exception¹², all non-managers we met are elected union representatives, either at CGT (*Confédération Générale du Travail*, left) or CFDT (*Confédération Française Démocratique du Travail*, centre-left). As such, they benefit from the allocation of part or all of their working time to the service of their union, without a loss in earnings (*délégation syndicale*). This *délégation* is part-time only for Nadia and Alain, but full-time for Youssef and Thérèse.

Companies A, B, C and D are of various sizes, but all have thousands of employees spread across hundreds or more of sites. The clients to which the sites belong are diverse: Nadia and Malika work in a university; Youssef also used to work there, and before also worked as a cleaner in a shopping centre; Alain works in a clinic, as did Thérèse before becoming a full-time union representative.

¹²That of Malika. However, we still met her through the mediation of her supervisor Nadia, a union representative, and they had recently participated in a strike together.

D Averages by group (main specification)




Figure 8: Total yearly earnings in the treated and the control group, in 2015 euros.




Figure 9: Days worked per year in the treated and the control group.




Figure 10: Hours per week in the treated and the control group.




Figure 11: Hourly wage (in 2015 euros) in the treated and the control group.




Figure 12: Share employed in the treated and the control group.




Figure 13: Share receiving unemployment benefits and no earnings in the treated and the control group.




Figure 14: Share receiving unemployment benefits in the treated and the control group.




Figure 15: Share with stable establishment in the treated and the control group.

E Main specification, detailed results




Figure 16: Main specification. Log earnings




Figure 17: Main specification. Earnings' components




Figure 18: Main specification. Proxies for unemployment




Figure 19: Main specification. Stable establishment dummy

F Heterogeneity analyses

F.1 Balanced panel




Figure 20: Balanced panel. Log earnings




Figure 21: Balanced panel. Earnings' components

F.2 Stayers only




Figure 22: Stayers at the same employer only. Log earnings




Figure 23: Stayers at the same employer only. Earnings' components

F.3 By destination industry

F.3.1 Catering and cleaning




Figure 24: Heterogeneity by industry (catering and cleaning). Log earnings




Figure 25: Heterogeneity by industry (catering and cleaning). Log days




Figure 26: Heterogeneity by industry (catering and cleaning). Log hours per day




Figure 27: Heterogeneity by industry (catering and cleaning). Log hourly wage




Figure 28: Heterogeneity by industry (catering and cleaning). Employment dummy




Figure 29: Heterogeneity by industry (catering and cleaning). Only unemployment benefits dummy




Figure 30: Heterogeneity by industry (catering and cleaning). Some unemployment benefits dummy




Figure 31: Heterogeneity by industry (catering and cleaning). Stable establishment dummy

F.3.2 Security and logistics




Figure 32: Heterogeneity by industry (security and logistics). Log earnings




Figure 33: Heterogeneity by industry (security and logistics). Log days




Figure 34: Heterogeneity by industry (security and logistics). Log hours per day




Figure 35: Heterogeneity by industry (security and logistics). Log hourly wage




Figure 36: Heterogeneity by industry (security and logistics). Employment dummy




Figure 37: Heterogeneity by industry (security and logistics). Only unemployment benefits dummy




Figure 38: Heterogeneity by industry (security and logistics). Some unemployment benefits dummy




Figure 39: Heterogeneity by industry (security and logistics). Stable establishment dummy

F.4 By sex




Figure 40: Heterogeneity by sex. Log days




Figure 41: Heterogeneity by sex. Log hours per day




Figure 42: Heterogeneity by sex. Log hourly wage




Figure 43: Heterogeneity by sex. Employment dummy




Figure 44: Heterogeneity by sex. Only unemployment benefits dummy




Figure 45: Heterogeneity by sex. Some unemployment benefits dummy




Figure 46: Heterogeneity by sex. Stable establishment dummy

F.5 By migration status




Figure 47: Heterogeneity by migration status. Log days




Figure 48: Heterogeneity by migration status. Log hours per day




Figure 49: Heterogeneity by migration status. Log hourly wage




Figure 50: Heterogeneity by migration status. Employment dummy

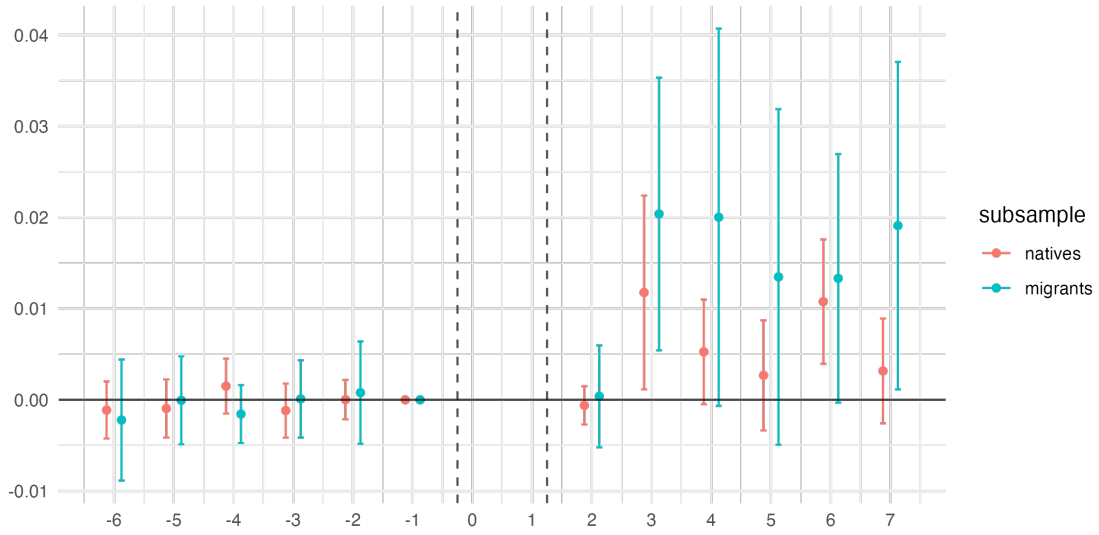


Figure 51: Heterogeneity by migration status. Unemployment benefits dummy




Figure 52: Heterogeneity by migration status. Stable establishment dummy

G Robustness checks

G.1 Including small transfers




Figure 53: Including small outsourcing transfers. Log earnings




Figure 54: Including small outsourcing transfers. Earnings' components




Figure 55: Including small outsourcing transfers. Proxies for unemployment




Figure 56: Including small outsourcing transfers. Stable establishment dummy

G.2 Without controls on dependent variables




Figure 57: Without controls on the dependent variables. Log earnings




Figure 58: Without controls on the dependent variables. Earnings' components




Figure 59: Without controls on the dependent variables. Proxies for unemployment




Figure 60: Without controls on the dependent variables. Stable establishment dummy

G.3 Without any controls




Figure 61: Without any controls. Log earnings




Figure 62: Without any controls. Earnings' components




Figure 63: Without any controls. Proxies for unemployment




Figure 64: Without any controls. Stable establishment dummy

G.4 $t - 2$ as reference




Figure 65: $t - 2$ as reference. Log earnings




Figure 66: $t - 2$ as reference. Earnings' components




Figure 67: $t - 2$ as reference. Proxies for unemployment




Figure 68: $t - 2$ as reference. Stable establishment dummy