Journées de méthodologie statistique de l'Insee 2022

La vie familiale explique-t-elle l'érosion de l'offre de travail des infirmières à l'hôpital ?

Pierre PORA (*)

(*) Drees, Crest et Université Paris Nanterre

pierre.pora@sante.gouv.fr

Mots-clés. : Soins infirmiers, offre de travail maternelle, event-study, différence-de-différences.

Domaines. 8.2 Économétrie appliquée, 8.3 Évaluation des politiques publiques.

Résumé

Cet article quantifie dans la contribution de l'effet causal de la maternité sur l'offre de travail des femmes au profil de cycle de vie des infirmières hospitalières françaises. L'insuffisance de l'offre d'infirmières est une préoccupation dans la plupart des économies avancées depuis plus de deux décennies (voir par exemple Shields, 2004). Le travail infirmier ayant un impact direct et positif sur la santé des patients (Propper and Van Reenen, 2010; Gruber and Kleiner, 2012) cette pénurie peut conduire à des effets néfastes sur la santé. Cette préoccupation est renforcée dans le contexte de l'épidémie de Covid-19, qui a encore accentué la pression portant sur les professionnels de santé, notamment à l'hôpital.

Le caractère très féminisé de la profession infirmière et des travaux récents montrant que l'offre d'infirmières peut réagir fortement aux politiques de congé parental (Friedrich and Hackmann, 2021) suggèrent que les difficultés à concilier vie familiale et vie professionnelle, et les normes de genre concernant l'éducation des enfants pourraient en fait expliquer une part substantielle de l'insuffisance de l'offre de travail des infirmières.

Je m'appuie sur des données administratives longitudinales détaillées issues à la fois des Déclarations Administratives de Données Sociales (DADS) et des registres de naissance chaînés au sein de l'Échantillon Démographique Permanent (EDP) pour explorer cette question. Je montre dans un premier temps que (i) l'offre de travail moyenne des infirmières hospitalières dans le secteur salarié diminue considérablement après leur premier emploi dans un hôpital, et (ii) le nombre de mères parmi elles augmente très fortement au cours des mêmes années.

Je mets ensuite en œuvre un cadre particulier de différence-de-différences appelé *event-study*, inspiré de Kleven, Landais, and Søgaard (2019) afin d'identifier l'effet causal des enfants sur l'offre de travail des mères. Je montre que la maternité conduit les infirmières à diminuer

Je remercie Magali Dumontet, Bertrand Garbinti, Paul Malliet, Éric Maurin, Dominique Meurs, Sarah Nedjar-Calvet, Roland Rathelot et Anne Solaz, ainsi que les participants à la conférence TEPP (Évry 2021) et aux JESF 2021, et aux séminaires de la Drees et de l'Université Paris Nanterre pour leurs commentaires et suggestions. Toutes les erreurs et opinions exprimées dans cet article sont de mon fait.

leur offre de travail dans le secteur salarié d'environ 0,15 équivalent temps plein au cours des dix premières années suivant la naissance de leur premier enfant. Cette diminution est entièrement due à des transitions vers des postes à temps partiel, par opposition à la décision de conserver ou pas un emploi salarié.

Enfin, j'utilise ces effets estimés pour comparer les profils de cycle de vie observés d'offre de travail des infirmières avec les profils contrefactuels qui seraient observés si l'effet des enfants était fixé à 0, c'est-à-dire (i) si les infirmières n'avaient pas d'enfants ou (ii) si les mères prenaient leurs décisions d'offre de travail de la même manière que les pères, dans la mesure où les pères ne réduisent pas leurs heures de travail à l'arrivée des enfants. Dans un tel cas, la décroissance des heures travaillées en emploi salarié au cours des dix premières années de la carrière serait 37% moins forte, et serait réduite de moitié en ce qui concerne les heures travaillées dans le secteur public.

Abstract

Soon after they land their first job at a hospital, many French nurses become mothers and decrease their hours worked in the salaried sector. I quantify the contribution of motherhood to the labor supply lifecycle profile of French hospital nurses. My event-study estimates, based on administrative registers, show that children cause female nurses to decrease their hours worked in the salaried sector by 0.15 full-time units in average. This decrease is entirely driven by the intensive margin of labor supply, as children do not induce nurses neither to leave salaried employment nor to turn to other jobs. Motherhood explains over a third of the drop in average hours worked in the salaried sector over the first ten years of a career, and half of the decline in nursing labor supplied to the public sector.

1 Introduction

The insufficiency in the supply of nurses has been a concern in most advanced economies for over two decades (see Shields, 2004, for instance). Because nursing labor has a direct and positive impact on patients' health outcomes (Propper and Van Reenen, 2010; Gruber and Kleiner, 2012; Friedrich and Hackmann, 2021), this shortage may be conducive to adverse health effects. The Covid-19 crisis, that further enhances the pressure on healthcare workers, makes this concern particularly salient, leading both policy makers and experts to worry about potential increases in outflows and decreases in inflows. In France specifically, these considerations triggered substantial pay increases for healthcare workers in the public sector in 2020-2021, that exceed +15% of their net wages, expecting these to discourage current workers to end their careers, and to attract additional workers. France ranks indeed quite low in the distribution of nurses' wages among OECD countries, and does not benefit from large foreign-born and foreign-trained nurses inflows (Lafortune, Socha-Dietrich, and Vickstrom, 2019).

The extent to which such pay increases are likely to restrain healthcare workers, and especially nurses, from leaving their occupation and to induce others to get into these occupations has been debated by researchers. Noting that the vast majority of nurses are women, a large share of this literature takes its roots in the female labor supply literature, in which broadly speaking the labor supply decision of nurses, at both margins, is assumed to depend on their (potential) wage rate, their non-earned income – including their partner's earnings – and constraints related to family life, e.g. whether they have young children (see Antonazzo et al., 2003, for a survey). Empirically, these papers tend to find positive albeit small own-wage elasticities, which implies that while these policies may indeed increase the labor supply of nurses, they will not be sufficient to achieve large surges. These small elasticities may arise if nurses are somewhat constrained by norms in their time allocation decisions. The vast majority of nurses being

women, and traditional norms that tend to restrict the labor supply of women make gender norms worth the investigation.

This paper approaches these issues by putting motherhood at the center of the analysis. Specifically, I quantify how much the causal effect of motherhood on women's labor supply explains of the labor supply lifecycle profile of French hospital nurses. Indeed, while several papers note that the presence of young children correlates with decreases in the labor supply of nurses – either in terms of hours worked (Askildsen, Baltagi, and Holmås, 2003; Hanel, Kalb, and Scott, 2014) or in terms of participation decision (Phillips, 1995; Nooney, Unruh, and Yore, 2010) –, very few of them further investigate this fact.[1] By contrast, I build on recent insights from the gender gap literature, according to which most of the differences in labor market outcomes between men and women result from multiple decisions, and especially negative labor supply decisions along multiple margins, that mothers make in response to the arrival of children (Kleven, Landais, and Søgaard, 2019). Nurses being an extremely feminized profession – in France, over 85% of nurses are women (Bessière, 2005) – and recent evidence that the supply of nursing can react strongly to parental leave policies (Friedrich and Hackmann, 2021) suggest that poor family-work conciliation and gender norms regarding child rearing may actually explain a substantial share of the insufficiency of nurses labor supply.

I leverage detailed longitudinal administrative data issued from both payroll tax forms and birth registers to explore this question. Such data enable me to track the salaried labor supply and the fertility decisions of a representative sample of qualified hospital healthcare workers, the vast majority of whom are nurses, over 30 years of their lifecycle, from 1988 to 2017. This allows me to show that (i) their average labor supply in the salaried sector decays substantially after they take their first job at a hospital, and (ii) the number of mothers among them increases very steeply over the same years.

I then implement an event-study framework in order to identify the causal effect of children on mothers' labor supply. This approach takes advantage from differences in the timing of the first childbirth across actual mothers of the same age who took their first job at as a qualified health worker at a hospital at the same time to identify the consequences of motherhood thanks to a limited anticipation and a parallel trends assumptions. Specifically, I investigate the effect of motherhood on labor supply decisions at multiple margins: participation, hours worked, occupation, work setting, sector. I show that motherhood induces nurses to decrease their labor supply in the salaried sector by about 0.15 full-time units during the first 10 years after the birth of their first child. This decrease is entirely driven by transitions to part-time positions, as opposed to participation decisions.

I finally use these estimated effects to compare the observed lifecycle profiles of nurses labor supply with counterfactual ones that would be observed if the effect of children is set to 0, that is either (i) if female nurses did not have children or (ii) if mothers made their labor supply decisions the same way as fathers do, since fathers do not reduce their hours worked upon the arrival of children. I find that in such a case, the decay in hours worked in salaried employment over the first ten years of a career would be 37% less steep, and would be cut by half when it comes to hours worked in the public sector.

The remainder of this paper is organized as follows. Next section outlines the institutional context. Section 3 describes the administrative data upon which the analysis is based. Section 4 details my empirical framework. Section 5 presents the results and Section 6 concludes.

[1] Other authors do not detect a substantial correlation between child-rearing and nurses labor supply decisions (Holmås, 2002; Estryn-Béhar et al., 2007; Frijters, Shields, and Price, 2007; Toren et al., 2012), or even find a negative correlation between motherhood and nurses' intention to quit (Shields and Ward, 2001).

2 Institutional context

Similar to what it the case in many countries, in France nursing is a licensed occupation. In other words, there a strong barriers to entry in the nursing market, as one has to meet several requirements to be granted the authorization to work as a registered nurse. The most salient of them is unsurprisingly education: authorization is only granted upon the completion of a curriculum at specific institutions (*Instituts de Formation en Soins Infirmiers*, IFSI). Getting into these institutions usually involves passing a competitive entrance exam, after which students follow a 3-years training program. Since 1979, the number of open positions in these competitive exams is fixed *ex ante* at the national level by the ministry of Health. These programs were part of vocational education up until 2009, but have since moved to higher education, and are nowadays provided by or in partnership with universities. As a result, this initial nurse training now corresponds to a bachelor degree.

This nursing degree grants with the authorization to work as a general care nurse. Over the course of their careers, and usually conditional on experience, nurses can choose to gain additional training to get into one of several nursing specializations. This additional training may either be part of a master degree or a professional degree, depending on the specialization at stake.

There are two main settings in which nurses may work in France. Firstly, nurses may work as salaried employees, whose employer may be a hospital, a long-term care facility, a health center, a school etc. Secondly, they may work as freelance nurses, in which case they provide healthcare directly to patients. However, freelance nursing is not open to fresh nursing graduates. Indeed, being granted the authorization to work as a freelance nurse requires not only a nursing diploma, but also at least two years of experience as a salaried nurse at a healthcare facility. In 2006, the vast majority of registered nurses (63%) were salaried employees at a hospital, either public (49%) or private (14%); 4% of them were employed in a long-term care facility, and 21% of them were salaried employees in other settings. Lastly, 12% of them worked as freelance nurses (Barlet and Cavillon, 2010). In this paper, I focus on the lifecycle of nurses who hold a job at a hospital at one point in their lives. This actually covers the vast majority of the nursing profession, because over three quarters of nurses work at a hospital when they begin their careers (see Appendix B).

Hospitals frequently offer daycare services to their staff, in order to make family-work conciliation easier for those jobs with long and atypical hours, and with frequent and impredictible changes of schedule (Daune-Richard, Odena, and Petrella, 2007). Unfortunately, quantitative data regarding these services do not seem available.

3 Data

My analysis is based on a combination of labor market data issued from payroll tax forms and fertility data issued from birth registers, all made available by Insee. I merge these datasets thanks to a common individual identifier based on a Social Security number. This allows me to build a sample of qualified healthcare workers who have at least once held a job at a French hospital, either in the public or the private sector, that I am able to track over time in the salaried sector from 1988 to 2018. Simple summary statistics based on this longitudinal data show that, soon after they land their first job at a hospital, (i) the average hours worked in the salaried sector by qualified healthcare workers starts to decay, and (ii) a large number of them become mothers during the same time-period.

3.1 Labor market data

My labor market data are drawn from the *Déclarations Annuelles de Données Sociales* (DADS). By law,[2] French employers have to fill in a DADS form for every employee subject to payroll taxes. The form contains detailed information about days paid, hours paid, occupation, industry, gross and net wages, other job characteristics (beginning, duration and end of a period of employment and part-time employment), employer characteristics (size and location) and individual characteristics (age, gender and municipality of residence). In Appendix A, I provide further details on how time worked is measured, and especially on how paid maternity leave is included in my measure of labor supply. Throughout the paper, my main variable of interest corresponds to hours worked, measured in full-time units. This value is capped to 1 for individuals working full-time during an entire year, so that it does not incorporate overtime.[3]

I take advantage of a longitudinal declination of these data. Specifically, I rely on the DADS panel, a longitudinal sample to track mothers' labor supply from 1988 to 2017, thanks to an anonymized personal identifier based on their social security number that allows me to link this information to birth records. The sampling rate of this longitudinal dataset varied over time: from 1988 to 2001, the data only cover individuals born in October in even-numbered years; as of 2002, it also includes individuals born on January 2-5, April 1-4, July 1-4 and October 1-4 regardless of their year of birth. This creates left-censoring regarding beginning of the career for the latter group of individuals. For this reason, I restrict the analysis to individuals who belong to the former group.

These data have two main caveats with respect to my analysis. The first one is that the most detailed occupation variable is not available before 2009. Indeed, before this employers only had to answer a 2-digits occupation question, as opposed to the 4-digits occupation which is the most detailed level in the occupation classification used by Insee. This prohibits the naive approach to the labor supply of nurses, which would basically select individuals into the sample based on whether or not they have, at one point in their lives, held a job as a nurse, as made salient by the 4-digits occupation variable. Instead, I select individuals based on the combination of the 2-digits occupation variable, and the 5-digits industry variable. Subsection 3.3 details this choice and characterizes the selected individuals in terms of their detailed occupation, when observed.

The second issue is that hours worked are not observed before year 1995. Before this, the data only provide information on days worked, and working-time status, either full-time or part-time. I choose to impute hours worked, measured in full-time units, before 1995, based on those two variables. Specifically, for full-time workers, I consider time worked to be proportional to days worked, and equal to 1 for those who work for a full year – in this there is absolutely no difference with the way time worked is measured after 1995. For part-time workers, I consider them to be on a 50% schedule, which was the most frequent case after 1995 (see Appendix A); as a result, time worked is proportional to days worked, so that for those who work for an entire year time worked is equal to 0.5. My results are nevertheless robust to this particular choice (see Subsection 5.3).

3.2 Fertility data

My analysis also relies on birth records. Births are registered by an individual who was present at the time of birth, usually the father, but in some cases a doctor or a midwife. I take advantage of a longitudinal version of these records at the individual level extracted from the *Échantillon Démographique Permanent* (permanent demographic sample, EDP) to obtain information on

[2] The absence of DADS as well as incorrect or missing answers are punished with fines.

[3] In Appendix C.3, I replicate my results, this time including overtime. I also explore changes in hourly wages, which are tightly linked to changes in working conditions, e.g. shift work. I find that the inclusion of overtime leads to very similar conclusions, and that hourly wages is left virtually unaffected by motherhood.

the timing of births. Because it displays an anonymized personal identifier based on the same social security number as the DADS data, this dataset can be merged with the longitudinal version of the DADS.

This dataset covers individuals born on October 1-4 whatever their year of birth; information regarding individuals born on January 2-5, April 1-4 and July 1-4 is available from 2004. To get around this left-censoring issue, and due to the sampling of the labor market data, I restrict my analysis to individuals born on October 1-4 of even-numbered years.

A caveat of this dataset is that some birth-related data for the 1990s were incomplete in administrative birth records for individuals born on October 2-3 (for details, see Wilner, 2016). For these individuals I use 1990 and 1999 census data to fill in the gaps, as do Pora and Wilner (2019). The quality of these data is comparable to that concerning individuals born on October 1 or 4 for whom administrative birth records are available from 1967.

3.3 Sample construction

My analysis relies on a sample of qualified healthcare workers working in a hospital setting, that I follow over the course of their lives. As explained in Subsection 3.1, data regarding the detailed occupation are not available before 2009, which restricts the possibility to base my sample selection on this variable. Instead, I rely on the combination of the 2-digits occupation and 5-digits industry. Specifically, I define hospital nurse jobs as those with (i) 2-digits occupation variable equal to "Qualified healthcare and social workers"[4] and (ii) 5-digits industry variable equal to "Hospital activities". This definition only approximates the usual definition of a hospital nurse job. However, it matches the usual definition quite closely: Table 1 displays the distribution of detailed occupations among jobs that match this criterion, which is observed in the DADS data as of 2009. Over three quarters of these jobs are indeed nurse jobs; the remainder are mostly health related technical jobs. Note that my approach includes nurses managers, but excludes auxiliary nurses who belong to another 2-digit occupation group. Social workers, although nominally included in the criterion, represent a very small share of this population.

Individuals of interest are all those that are observed, between 1988 and 2017, to hold this type of job for at least six months. I then track their labor market trajectories from 1988 to 2017, regardless of whether or not they still hold this kind of job. In other words, "nurses" who form part of my sample are not necessarily nurses throughout their lives, but it is assured that they have been at some point.[5] This leaves me with 161,723 observations that account for 5,627 individuals. As detailed above, inclusion in this sample is based on individual's birthday; as long as the causal effect of birthday on labor supply is close to 0, it is therefore a representative sample of the population of interest at rate 0.6%.

The data only allow to follow individuals into salaried employment. When individuals are not observed with a salaried job, I consider their hours worked to be equal to 0. In that way, my estimates regarding regarding hours worked are not conditional on salaried employment, as they do take into account time-periods spent outside employment, but they do not take into account the labor supply in the freelance sector.

Crucial to my analysis are (i) the year during which individuals are first observed to hold a nurse job at a hospital and (ii) the length of the time-period during which they are observed afterwards. That my data only cover the 1988-2017 time-period generates two issues with respect to this. Firstly, a substantial share of individuals who are observed to hold such a job in 1988 are

[4]Medical doctors do not belong to this category, as they belong to another one that is part of the "Managers and professionals" group.

[5]Additionally, the vast majority of individuals who ever hold a nurse position will fall within the universe that I cover, because most of them work at a hospital at some point of their lives. To show this, in Appendix B I quantify the share of workers who ever hold a job at a hospital among those who I can properly identify as beginning their careers as nurses after 2010 thanks to the detailed occupation variable. This share is about 75%.

likely to have done so for an unobserved number of years, which creates a left-censoring issue. Secondly, recent cohorts, as defined based on the timing of the first hospital nurse job, are only observed for a restricted time-period afterwards. My identification strategy is entirely based on within-cohort comparisons of nurses who become mothers sooner or later. Hence, for these recent cohorts, some control groups of mothers whose first child is born later are unobserved, because the birth of this child cannot be recovered from the data. This creates a right-censoring issue. I discuss these two concerns in Subsection 5.3 and provide evidence that my results are robust with respect to these issues.

3.4 Summary statistics

Table 2 displays a few summary statistics regarding the sample. Individuals of interest usually get their first job as a hospital nurse around age 30, usually after a few years in different jobs, either as a nurse in a different work setting, or in a different occupation. By age 45, about 80% of them have children. This rate is comparable to that of the overall French population (81,7% for women, see Reynaud, 2020). As in 2015 in France, mothers' average age when they gave birth to their first child was 28.5, this suggests that for a large share of hospital nurses, the beginning of their career coincides with child-rearing years. The remainder of the paper aims at quantifying this fact and its implications.

To this end, Figure 1 plots the share of mothers over time relative to the first hospital nurse job, both in the aggregate and separately across cohorts defined by the timing of the first hospital nurse job. The share of mothers increases slightly before nurses land their first job at a hospital, but remains quite small: one year before they get their first job as a qualified healthcare worker at a hospital, less than one in five nurses are mothers. By contrast, after they get their first nurse job at a hospital, this share rises quickly: ten years after the said first nurse job at a hospital, the share of mothers is 66%. Note that this share is not conditional on gender: among women the proportion is about 80%. For the sake of this particular paper, considering the motherhood rate unconditional on gender has merits because what matters for the aggregate supply of nursing labor is the unconditional rate, as opposed to the fertility rate of women. The difference between the rate at the onset of a career and the rate several years afterwards implies a quick expansion of the number of mothers soon after nurses land their first job at a hospital, as made obvious by the figure. Splitting the data across cohorts confirms that this pattern is not driven by changes in the composition of cohorts that are observed at each point of time, given the restricted time-period of observation.

Figure 2 displays the labor supply, that is (a) the average hours worked, measured in full-time units and (b) the salaried employment rate, over time relative to the first hospital nurse job, both in the aggregate and across cohorts defined by the timing of the first hospital nurse job. Soon after the first two years, which correspond to a gradual entry into the job, the labor supply in the salaried sector starts to decay. This decline amounts to 0.16 full-time units over the first ten years of a career, which is substantial considering that the baseline at the beginning of a career is about 0.85 full-time unit. This decline is to a large extent driven by decisions at the extensive margin, i.e. participation decisions, as made obvious by the large decrease in the salaried employment rate. Once again, this pattern holds across cohorts and is not driven by changes in composition over time.

The data allow me to further details these results. Indeed, the DADS data are available at the job spell level, that is at the individual $\times$ employer $\times$ year. As a result, hours worked by an individual at a given point in time can be decomposed as the sum of hours worked across all employers. Specifically, in this paper, I consider three decompositions. Firstly, I contrast hours worked as a nurse with hours worked in other jobs: nurse jobs are those that belong to the "Qualified healthcare and social workers" 2-digit occupational group that I use in the sample definition. Secondly, I compare hours worked at a hospital with hours worked in other work

settings: hours worked at hospital is the sum of hours worked across employers that belong to the "Hospital activities" 5-digit industry according to the data. Lastly, I confront hours worked in the public sector with hours worked in the private sector, based on the sector to which employers belong. Because inclusion in my sample is based on whether individuals hold a specific type of job once in their lives, these decompositions are useful to document whether hospital nurses transition to other types of jobs over the course of their lifecycle.

Figure 3 relies on these decompositions to compare the lifeycle profiles of (a) hours worked at a hospital vs. in another setting; (b) hours worked as a nurse vs. in another occupation; and (c) hours worked in the public sector vs. in the private sector. It makes it clear that the decline in labor supply in the salaried sector is first and foremost driven by a decline in hours worked at a hospital, as a nurse, and to a lesser extent in the public sector. Interestingly, although hours worked in non-hospital or non-nurse jobs do slightly increase, the magnitude of this rise remains limited. This implies that the decay in hours worked as a hospital nurse over the lifecycle is not driven by the reallocation of hospital nurses towards other salaried jobs.

4 Empirical analysis

My analysis builds on the event-study approach proposed by Kleven, Landais, and Søgaard (2019). It slightly improves on it by: (i) using more restrictive comparison groups and (ii) incorporating insights from the recent difference-in-difference literature (see Callaway and SantAnna, 2020; de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2020; Goodman-Bacon, 2021). Specifically, my approach aims at preventing identification issues related to the use of two-way fixed effects in settings where treatment effects are likely to be heterogeneous. The exposition of my empirical framework is largely based on Callaway and SantAnna (2020) and Sun and Abraham (2020).

4.1 Model and identification

Let $Y_{i,t}$ denote the labor supply – i.e. the total number of hours worked in the salaried sector – of individual i at time t, which is measured relative to when she took her first job as a qualified health worker at a hospital. Let G_i denote the group to which individual i belongs, which is defined by (i) her year of birth and (ii) the year during which she first took a nurse job at a hospital.[6] Lastly let C_i denote the year during which her first child was born ($C_i = \infty$ if she is without child).

I define $Y_{i,t}(c)$ to be the potential labor supply of individual i at time t had she gave birth to her first child at time c. Consistently, $Y_{i,t}(\infty)$ is her labor supply at time t had she chosen to remain childless. By construction:

$$Y_{i,t} = Y_{i,t}(\infty) + \sum_{c} (Y_{i,t}(c) - Y_{i,t}(\infty)) \mathbb{1}\{C_i = c\} \tag{1}$$

My analysis revolves around the causal effect of motherhood on labor supply. In other words, I am interested in (functionals of) the distribution of random variables $Y_{i,t}(c) - Y_{i,t}(\infty)$, with $c < \infty$. Specifically, I define the cohort-specific average treatment effect on the treated:

$$CATT_{g,c,t} = \mathbb{E}[Y_{i,t}(c) - Y_{i,t}(\infty) \,|\, G_i = g,\, C_i = c] \tag{2}$$

This quantity corresponds to the effect of being $t - c$ years away from the birth of one's first child, for those who gave birth to their first child at time c, and belong to group g. These

[6] Due to the left-censoring if the data, this year is not observed for individuals who got their first hospital nurse job before 1989. These individuals are grouped together in groups defined by (i) the year of birth and (ii) having taken the first hospital nurse job in 1988 or before.

average treatment effects are not conditional on possible subsequent childbirths. As a result, it incorporates both the causal effect of motherhood at the extensive margin, i.e. choosing to be a mother or not, and that of to the intensive margin, i.e. choosing to give birth to one additional child for those who are already with child. In other words, the causal effect of motherhood mixes that of the first child, and of all subsequent children, with weights that depend on the difference between the time-period t and the timing of the first child's birth c: short-run effects ($t = c$) relate almost exclusively to the extensive margin of fertility, whereas longer run effects ($t > c$) will integrate a larger share of the consequences of the intensive margin. This is especially true in a context in which most parents choose to have more than one child, as implied by Table 2. I discuss these concern, and provide a decomposition of the effect of children between these two margins in Appendix C.2.1.

To identify these quantities from the data, I make two assumptions: (i) a parallel trend assumption and (ii) a limited anticipation assumption.

Assumption 1 (Parallel trends in baseline outcome)**.** *For all g, for all (t, t'), for all (c, c'), if $c, c' > 1$ and $c, c' < \infty$ then:*

$$\mathbb{E}[Y_{i,t}(\infty) - Y_{i,t'}(\infty) \,|\, G_i = g,\, C_i = c] = \mathbb{E}[Y_{i,t}(\infty) - Y_{i,t'}(\infty) \,|\, G_i = g,\, C_i = c'] \tag{3}$$

Assumption 2 (Limited anticipation)**.** *For all t, for all g, for all c, if $t < c - 1$ then:*

$$\mathbb{E}[Y_{i,t}(c) - Y_{i,t}(\infty) \,|\, G_i = g,\, C_i = c] = 0 \tag{4}$$

Assumption 1 states that absent children, the average labor supply of mothers who were born at the same time, took their first job at a hospital during the same year and had their first child only afterwards, would evolve in parallel over time. Assumption 2 states that the average effect of children on their mothers' labor supply is 0 up until two years before they are born. The reason for this choice, as opposed to a full no-anticipation assumption is that (i) becoming a mother during year t generally results from fertility decisions that were made during year $t-1$, and (ii) maternity leave may start during the last year before childbirth in case childbirth happens at the beginning of the civil year, which will mechanically affect the mother's labor supply.

Under these assumptions, provided that, within group, there is sufficient variation in the timing of childbirth, cohort-specific ATTs can be identified from the data.

Proposition 1 (Difference-in-difference estimand)**.** *For all (g, c, t), if $1 < c < \infty$ then:*

$$\begin{aligned} CATT_{g,c,t} \;=\; & \mathbb{E}[Y_{i,t} \,|\, G_i = g,\, C_i = c] \\ & -\mathbb{E}[Y_{i,c-2} \,|\, G_i = g,\, C_i = c] \\ & -\mathbb{E}[Y_{i,t} \,|\, G_i = g,\, \max(1,\, c-2\,, t+1) < C_i < \infty] \\ & +\mathbb{E}[Y_{i,c-2} \,|\, G_i = g,\, \max(1,\, c-2,\, t+1) < C_i < \infty] \end{aligned} \tag{5}$$

Proposition 1 implies that as long as, within a group, there are mothers to be that can still be observed at least two years before their first child is born, it is actually possible to impute the counterfactual labor supply lifecycle profile of mothers whose first child is already born, so as to identify cohort-specific ATTs. Specifically, let $\{\underline{T(g)},\, \underline{T(g)}+1,\, ...,\, \overline{T(g)}-1,\, \overline{T(g)}\}$ denote the set of time-periods that can be observed for individuals who belong to group g. Then for all c, $CATT(g, c, t)$ is identified from the data provided that:

(i) $\underline{T(g)} \leq c - 2 \leq \overline{T(g)}$;

(ii) $\underline{T(g)} \leq t \leq \overline{T(g)}$;

(iii) $\mathbb{P}(\max(1,\, c-2,\, t+1) < C_i < \infty \,|\, G_i = g) > 0$.

This last condition implies that very long run effects generally cannot be identified under these assumptions, because no counterfactual is available after the last mother is about to have her first child. Specifically, given the profile of Figure 1, cohort-specific effects are very unlikely to be identified for $t \geq 10$, and even more so for $t \geq 15$.

4.2 Aggregation and estimation

Aggregation The quantities I am interested in correspond to the causal effect of having been a mother for a certain amount of time, for women who, before becoming mothers, held a job as a qualified healthcare worker at a hospital. I recover these quantities by aggregating my cohort-specific ATTs with weights proportional to sample size. Specifically, let Ω be the subset of group-cohort-time-period triplets for which all three conditions hold, as well as $C_i > 1$, so that $CATT(g, c, t)$ is identified from the data. I define:

$$\tau(s) = \mathbb{E}[Y_{i,C_i+s}(C_i) - Y_{i,C_i+s}(\infty) \,|\, (G_i, C_i, C_i + s) \in \Omega] \tag{6}$$

This quantity represents the average treatment effect of being s years away from the birth of one's first child, for a certain subset of individuals, which varies depending on s. By the law of iterated expectations:

$$\tau(s) = \sum_{(g,c,c+s)\in\Omega} \mathbb{P}(G_i = g, C_i = c \,|\, C_i(G_i, C_i, C_i + s) \in \Omega)\, CATT(g, c, c + s) \tag{7}$$

By Proposition 1, it is therefore possible to express $\tau(s)$ as a function of quantities that are all identified from the data.

Lastly, I consider $\bar{\tau}(S)$, a quantity that represents the impact of children over the first S years of motherhood:

$$\bar{\tau}(S) = \frac{1}{s} \sum_{s=0}^{S} \tau(s) \tag{8}$$

Estimation Combined with Proposition 1, Equation 7 suggests a very simple plug-in estimator, in which population probabilities and expectations are replaced by their empirical analogues. The same also goes for the estimation of $\bar{\tau}(S)$.

Under usual integrability assumptions, these estimators are asymptotically normal (Callaway and SantAnna, 2020). To conduct inference, I choose to rely on a bootstrap approach, clustered at the individual level. This level of clustering is justified both from a sampling perspective – as the sampling scheme is defined at the individual level – and from a design perspective – as the treatment, i.e. children, is assigned at the individual level (Abadie et al., 2017).

4.3 Simulation exercise

To quantify the contribution of children to the lifecycle profile of nurses' labor supply, I build a counterfactual profile of nurses' labor supply, if either (i) nurses did not have children; or (ii) motherhood had no impact on the labor supply of nurses; or (iii) female nurses made their labor supply decisions the same way as men do, given that men usually do not decrease their labor supply upon becoming fathers (see Kleven, Landais, and Søgaard, 2019).

To this end, I first consider the realized lifecycle profile of average labor supply, described by $\mathbb{E}[Y_{i,t} \,|\, \underline{T(G_i)} \leq t \leq \overline{T(G_i)}]$, which corresponds to the average hours worked by nurses observed t years after they get their first job at a hospital. I consider the counterfactual lifecycle profile of labor supply to be described by the quantity:

$$\Lambda(t) = \mathbb{E}[Y_{it} \,|\, \underline{T(g)} \leq t \leq \overline{T(g)}] - \sum_{s\geq 0} \mathbb{P}(t = C_i + s, C_i > 1 \,|\, \underline{T(G_i)} \leq t \leq \overline{T(G_i)})\, \tau(s) \tag{9}$$

This approach: (i) only focuses on children born *after* mothers get their first job as a qualified healthcare worker at a hospital; and (ii) abstracts from considerations related to treatment effect heterogeneity across cohorts, and related compositional shifts.

5 Results

5.1 Children-related labor supply decisions

Figure 4 displays the results of my event-study framework, that is average changes in hours worked by mothers whose first child is s years old, relative to (i) 2 years before this child was born, and (ii) the change in labor supplied by mothers to be over the same time-period. Specifically, I consider (a) hours worked and (b) the salaried employment rate. Under a parallel trend and a limited anticipation assumption, these quantities identify the causal impact of motherhood on mothers' labor supply, that is the average treatment effect for a certain subpopulation (see Section 4).

First, before they have children, the dynamics of female nurses' labor supply mimick those of their counterparts who are to become mothers at a different time. Specifically, the differences in the average change in both hours worked in the salaried sector, measured in full-time units, and the salaried employment rate, across groups of mothers that are to have their children at different dates is a precisely estimated 0 over up to 10 years before the first childbirth. While this is not sufficient to assess the validity of my identifying assumptions in their full generality, it does support their credibility, as they impose the corresponding causal effect to be 0.

Second, after they have children, the dynamics of mothers' labor supply diverge depending on the timing of their first child's birth. Specifically, after the birth of their first child, mothers decrease their hours worked relative to mothers who will have their first child later on. This decrease corresponds both to the causal effect of the extensive margin of fertility, i.e. the decision to become a mother, and to that of the intensive margin of fertility, i.e. the decision to have additional children after the first one is born. Specifically, the short-run effects ($s = 0$) correspond mostly to that of the extensive margin. As time goes by ($s > 0$), these estimates put more and more weight on the intensive margin, given that most mothers choose to have additional children, as shown in Table 2.

Overall, over the first 10 years after their first child's birth, the magnitude of the decrease in hours worked induced by motherhood is about 0.15 full-time units. By contrast, the salaried employment rate of nurses does not decline due to motherhood. In other words, having children does induce nurses to diminish their hours worked, presumably by shifting to part-time positions, but does not lead them to leave the salaried labor force.

Figure 5 further decomposes the consequences of motherhood by replicating the analysis for different work-setting specific measures of labor supply. Specifically, I consider as my outcome hours worked in the salaried sector, measured in full-time units, contrasting: (a) hospital and non-hospital jobs; (b) nurse jobs and others jobs; and (c) public sector and private sector jobs. The figure makes it very clear that the impact of motherhood concentrates on public hospital nurse jobs: the effect on non-hospital, non-nurse or private sector jobs is close to 0. Part of the reason for this is that hours worked is a non-negative quantity, and that the baseline is not high (see Figure 3), which mechanically constrains the magnitude of potential drops.[7] However, the key message is that hospital nurses do not reallocate to other, plausibly more family-friendly

[7] More generally these comparisons are not conditional on having positive labor supply in one or the other sector, so that they do not contrast *the intensity* of the child penalty conditional on being e.g. a public sector nurse vs. a private sector nurse. They are informative as to where the aggregate drop in labor supply comes from, which mixes up (i) the fact that nurses are more numerous in the public sector than they are in the private sector, and (ii) the fact that conditional on sector choice, the impact of motherhood may differ across sectors.

job, upon becoming mothers: this would imply positive effects on hours worked in non-nurses, non-hospital jobs, which are rejected on the data.

In Appendix C.2, I delve further into the data to expand these results. The main lessons are that (i) these aggregate results are presumably driven by the extensive margin of fertility decisions; (ii) these labor supply decisions result in substantial earnings drop; (iii) motherhood does not seem to affect overtime or working conditions, as this would lead to decreases in hourly wages that are not observed; and (iv) men do not seem to decrease their labor supply upon becoming fathers, even though the related estimates are very imprecise due to the small sample size.

5.2 Contribution to lifecycle profiles of labor supply

I quantify the contribution of these children-related labor supply decisions to the decline of nurses' labor supply over the course of their career by comparing the observed profiles to counterfactual ones, if female nurses did not decrease their labor supply upon becoming mothers. In practice, I simply substract the causal effect of children, weighted by the share of mothers, from the observed profiles (see Subsection 4.3).

Figure 6 displays the results of this simulation exercise when it comes to (a) hours worked in the salaried sector, measured in full-time units, and (b) the salaried employment rate. As for hours worked, the dashed line, which represents the counterfactual lifecycle profile, progressively diverges from the plain line which corresponds to the observed profile, similar to that of Figure 2. However, shutting down mothers' labor supply decisions is not sufficient to rub away the decay in hours worked over the course of a career: nurses' labor supply directed towards the salaried sector would still decline even if either (i) nurses did not have children or (ii) female nurses made the same decisions as male nurses who do not reduce their hours worked upon becoming fathers. Specifically, in the simulated counterfactual, average hours worked would decrease by about 0.10 full-time units over the first 10 years of a career, against an observed 0.16. In other words, children explain slightly more than a third of the decline in hours worked over the 10 first years of a career. By contrast, and consistent with Figure 4, the salaried employment rate of nurses would stay the same even if nurses did not have children.

Figure 7 further details these results by replicating this comparison across different work settings. Once again, I contrast (a) hours worked inside or outside a hospital; (b) hours worked as a nurse or not; and (c) hours worked in the public and private sector. Because the effect of children only transits through hours worked at a hospital, as a nurse or in the public sector, the counterfactual and observed profiles are the same outside these work settings. Interestingly, the magnitude of the decline would be about 0.04 full-time units in the public sector absent motherhood, against a realized 0.1: children-related labor supply decisions thus explain more than half of the decay of hours worked in the public sector.

Appendix C.4 expands these results by considering the extensive margin of labor supply across work settings. Children contribute very little to the decline in nurses' salaried employment rate, except perhaps in the public sector where motherhood explains about 20% of the decline in the probability to hold a job after 10 years.

5.3 Robustness checks

Three main issues arise regarding the data upon which this analysis relies. The first two deal with restrictions on the time-period of observation; the last one with the measurement of hours worked in the DADS data.

Left-censoring In the DADS data, the beginning of the hospital nurse career is not observed for those who began their career in 1988 or before. This issue is made salient by the fact that

the parallel trends assumption (Assumption 1) is only made conditional on both the year of birth and the timing of the beginning of the career. In other words, mothers' counterfactual labor supply if they did not have children is only imputed based on women who were born at the same time and started their career the same year as them. In practice, I choose to gather all individuals who started their career in 1988 or before in one single group, so that among them the parallel trend assumption is only conditional on the year of birth. Figures D.1 and D.2 assess the robustness of my results with respect to this choice by replicating Figures 4 and 6 while omitting the data on these individuals. The results look extremely similar, confirming that they are immune to this particular issue.

Right-censoring The DADS-EDP dataset does not report childbirth that occurred after 2017. As a result, part of my control groups of mothers who are to have children later are not observed, as they cannot be differentiated from women who remain childless. This impedes the identification of cohort-specific ATTs for the year 2017, and more generally changes the composition of the control groups as one looks at youngers individuals. I check that my results are not affected by this issue by replicating Figures 4 and 6 while restricting to individuals for which the distinction between mothers who have children later and childless women is less problematic. Specifically, I consider this to be the case for individuals who got their first job as a hospital nurse before 2003, as they are likely to have completed their fertility decisions by 2017, at least 15 years after they became a hospital nurse (see Figure 1). Figures D.3 and D.4 display my results. They are very close to my baseline results.

Hours worked measurement Hours worked are not observed before 1995 in the DADS data. Before this, I impute hours worked based on days worked and working-time status (full-time or part-time), under the assumption that part-time workers are on a 50% schedule. This schedule is the most frequent among part-time workers when hours are observed (see Appendix A). To assess the robustness of my results with respect to this imputation, I replicate my analysis only while restricting to the 1995-2017 time-period. This will restrict the set of cohort-specific ATTs that can be identified under my identifying assumptions. Figures D.5 and D.6 display my results. They are once again very similar to my baseline results, which suggest that my baseline results are not driven by my choice of imputation.

Stable control group The event-study analysis upon which these results are based uses most individuals as both treated and control units: until two years before they give birth to their first child, all nurses belong to the control group, after which they become part of the treated group. As a result, the control group changes over time relative to the first child's birth: whereas, within a cohort, the short-run effect of motherhood relies on a control group that gathers almost all nurses but the one who become mothers immediately after they get their first job, the long-run effect relies on a control group that is restricted to nurses who become mothers long after they get their first job. To assess whether the dynamics of the treatment effects is driven by this compositional shift in the control group, I replicate my analysis, this time restricting the control group to nurses who have their first child at least 14 years after they got their first nurse job at a hospital, and restricting the treated group to nurses who have their first child at most 13 years after their first job. Additionally, I restrict the data to the 13 years that follow the first hospital nurse job. This approach is akin to more traditional difference-in-difference approaches in which units do not switch from one group to the other over time. Figure D.7 displays the resulting event-study estimates. They are very close to my baseline estimates, especially regarding hours worked. The employment effects appear slightly larger but are still compatible with my baseline results. Furthermore, Figure D.8 shows that even so, the contribution of motherhood to the labor supply lifecycle at the extensive margin remains marginal at best.

6 Conclusion

In this paper, I quantify the contribution of children to the decrease in hours worked by French hospital nurses in the salaried sector in the beginning of their careers. Building on an event-study framework and longitudinal administrative records, I find that motherhood causes nurses to decrease their average hours worked by about 0.15 full-time units after their first child is born. This decrease is almost entirely driven by changes in hours worked as a nurse, in a hospital and in the public sector, as opposed to hours in other occupations or work settings. Because the vast majority of nurses are women, and because many of them become mothers upon the beginning of their careers, motherhood affects the aggregate labor supply of nurses: overall, the impact of children on their mother's labor supply explains more than a third of the average decrease in hours worked in the salaried sector over the first ten years of a career, and about half of the decline in the public sector.

These findings imply that the supply of nursing labor should be studied in relation to fertility decisions, gender norms and family policies, as suggested by Friedrich and Hackmann (2021). Specifically, they suggest that the aggregate supply of nursing labor should increase as a result of either more women choosing to remain childless, or thanks to policies that further enhance family-work conciliation. Interestingly, while mothers do decrease their hours worked as a hospital nurse upon giving birth to their first child, my estimates do not suggest that they reallocate to other, more family-friendly, jobs or work settings, at least as far as the salaried sector is concerned. They also do not suggest that nurses are actually prone to retrieve from the salaried workforce upon becoming mothers. A remaining possibility is that the decrease in hours worked in the salaried sector is partially compensated by an increase in labor supplied to the freelance sector. While my results are not compatible with transitions at the extensive margin, additional data are required to delve further into this matter.

References

Abadie, A., S. Athey, G.W. Imbens, and J. Wooldridge. 2017. "When Should You Adjust Standard Errors for Clustering?" Working Paper No. 24003, National Bureau of Economic Research.

Antonazzo, E., A. Scott, D. Skatun, and R.F. Elliott. 2003. "The labour market for nursing: a review of the labour supply literature." *Health Economics* 12:465–478.

Askildsen, J.E., B.H. Baltagi, and T.H. Holmås. 2003. "Wage policy in the health care sector: a panel data analysis of nurses' labour supply." *Health Economics* 12:705–719.

Barlet, M., and M. Cavillon. 2010. "La profession infirmière: situation démographique et trajectoires professionnelles." Études et Résultats No. 759, Drees.

Bessière, S. 2005. "La féminisation des professions de santé en France: données de cadrage." *Revue française des affaires sociales* 1:17–33.

Callaway, B., and P.H. SantAnna. 2020. "Difference-in-Differences with multiple time periods." *Journal of Econometrics*.

Daune-Richard, A.M., S. Odena, and F. Petrella. 2007. "Entreprises et modes daccueil de la petite enfance. Innovation et diversification." Dossier d'étude No. 91, Cnaf.

de Chaisemartin, C., and X. D'Haultfœuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review* 110:2964–96.

Estryn-Béhar, M., B.I.J.M.V. der Heijden, H. Ogiska, D. Camerino, O.L. Nézet, P.M. Conway, C. Fry, and H.M. Hasselhorn. 2007. "The Impact of Social Work Environment, Teamwork Characteristics, Burnout, and Personal Factors upon Intent to Leave among European Nurses." *Medical Care* 45:939–950.

Friedrich, B.U., and M.B. Hackmann. 2021. "The Returns to Nursing: Evidence from a Parental-Leave Program." *The Review of Economic Studies*.

Frijters, P., M.A. Shields, and S.W. Price. 2007. "Investigating the quitting decision of nurses: panel data evidence from the british national health service." *Health Economics* 16:57–73.

Goodman-Bacon, A. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics*.

Gruber, J., and S.A. Kleiner. 2012. "Do Strikes Kill? Evidence from New York State." *American Economic Journal: Economic Policy* 4:127–57.

Hanel, B., G. Kalb, and A. Scott. 2014. "Nurses labour supply elasticities: The importance of accounting for extensive margins." *Journal of Health Economics* 33:94–112.

Holmås, T.H. 2002. "Keeping nurses at work: a duration analysis." *Health Economics* 11:493–503.

Kleven, H., C. Landais, and J.E. Søgaard. 2019. "Children and Gender Inequality: Evidence from Denmark." *American Economic Journal: Applied Economics* 11:181–209.

Lafortune, G., K. Socha-Dietrich, and E. Vickstrom. 2019. "Recent trends in international mobility of doctors and nurses." In OECD, ed. *Recent Trends in International Migration of Doctors, Nurses and Medical Students*. OECD Publishing, pp. 11–34.

Nooney, J.G., L. Unruh, and M.M. Yore. 2010. "Should I stay or should I go? Career change and labor force separation among registered nurses in the U.S." *Social Science & Medicine* 70:1874–1881.

Phillips, V.L. 1995. "Nurses' labor supply: Participation, hours of work, and discontinuities in the supply function." *Journal of Health Economics* 14:567–582.

Pora, P., and L. Wilner. 2019. "Child Penalties and Financial Incentives: Exploiting Variation along the Wage Distribution." Documents de Travail de l'Insee - INSEE Working Papers No. G2019/08, Insee.

Propper, C., and J. Van Reenen. 2010. "Can Pay Regulation Kill? Panel Data Evidence on the Effect of Labor Markets on Hospital Performance." *Journal of Political Economy* 118:222–273.

Reynaud, D. 2020. "Les femmes les plus modestes et les plus aisées ont le plus d'enfants." Insee

Première No. 1826, Insee.
Shields, M.A. 2004. "Addressing nurse shortages: what can policy makers learn from the econometric evidence on nurse labour supply?*." *The Economic Journal* 114:F464–F498.
Shields, M.A., and M. Ward. 2001. "Improving nurse retention in the National Health Service in England: the impact of job satisfaction on intentions to quit." *Journal of Health Economics* 20:677–701.
Sun, L., and S. Abraham. 2020. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics*.
Toren, O., R. Zelker, M. Lipschuetz, S. Riba, S. Reicher, and N. Nirel. 2012. "Turnover of registered nurses in Israel: Characteristics and predictors." *Health Policy* 105:203–213.
Wilner, L. 2016. "Worker-firm matching and the parenthood pay gap: Evidence from linked employer-employee data." *Journal of Population Economics* 29:991–1023.

Figures

Figure 1: Lifecycle profile of fertility: share of mothers among nurses of both genders

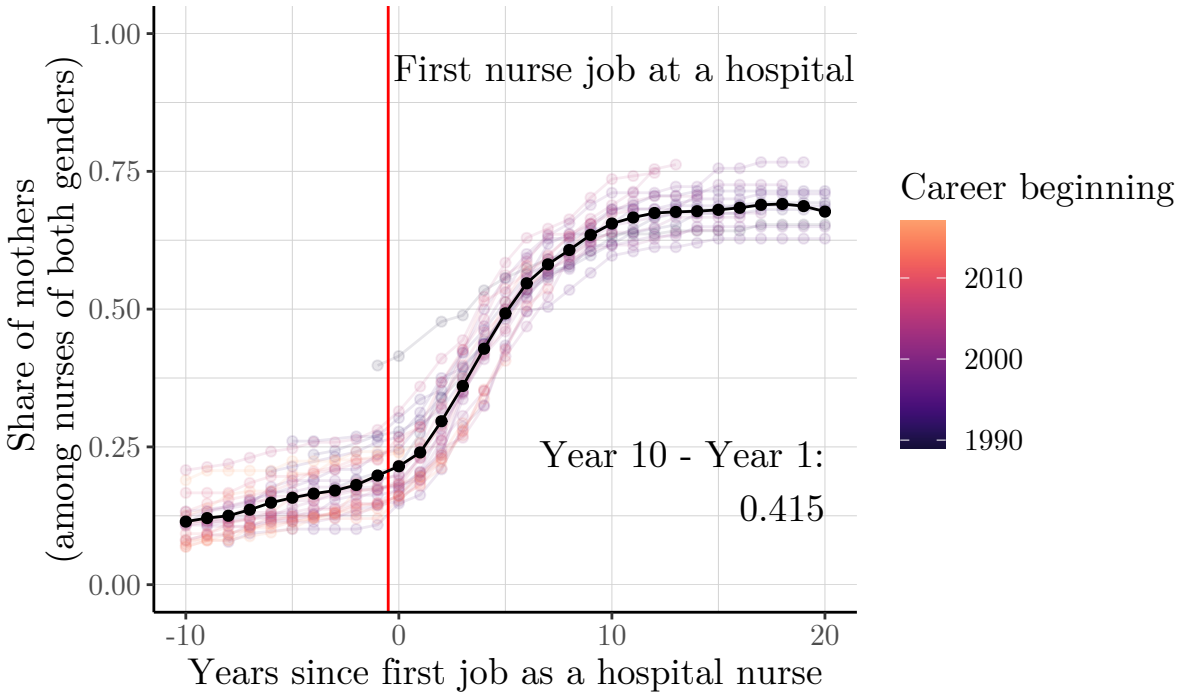

Share of mothers, by time relative to the first qualified healthcare worker job at a hospital.

Note. Data on individuals who got their first hospital nurse job in 1988 or before are omitted from the computation.

Source. Insee, DADS-EDP panel.

Figure 2: Lifecycle profile of hospital nurses' labor supply: total labor supply in the salaried sector

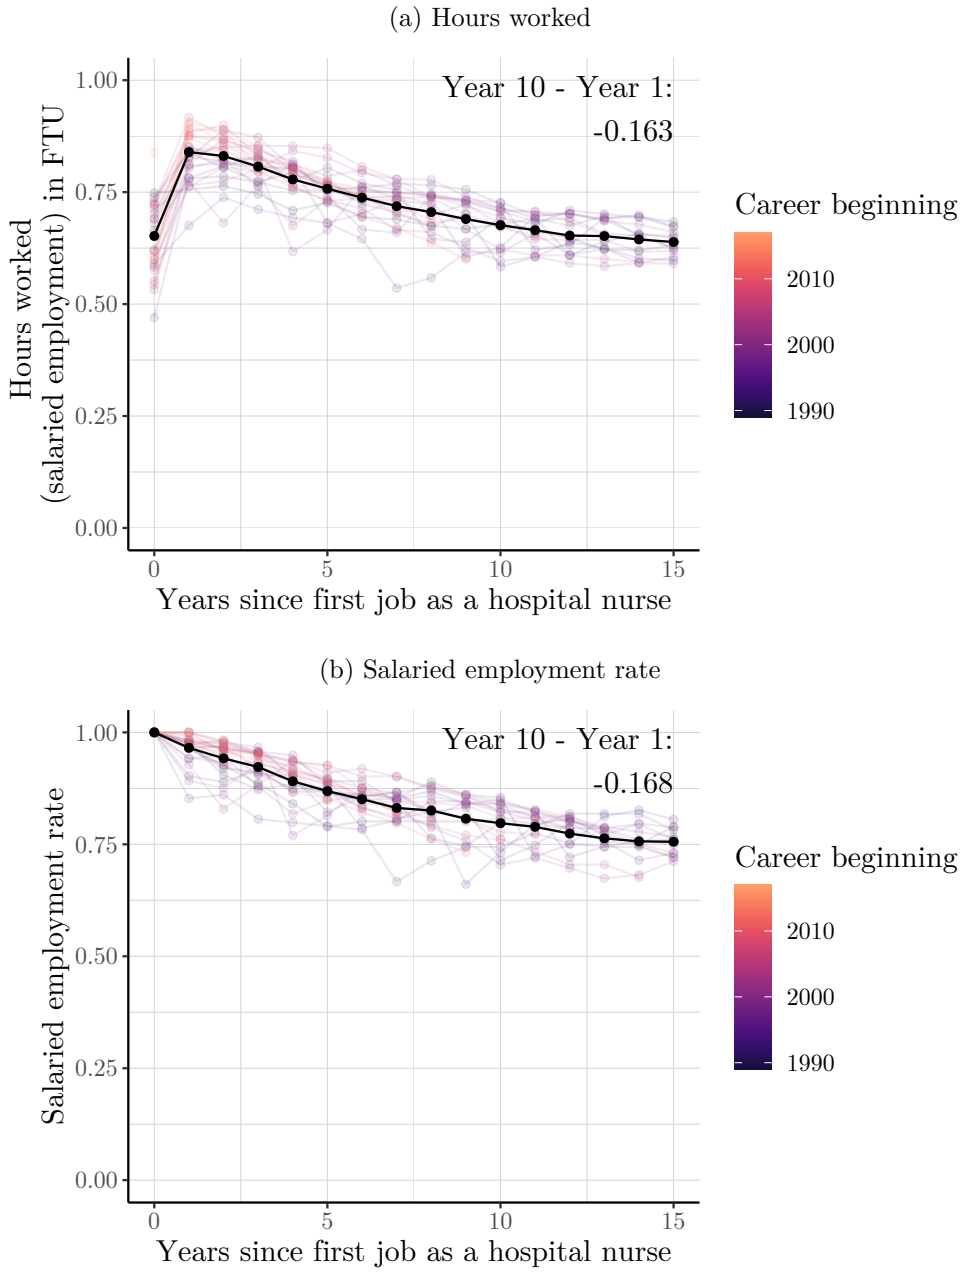

Average hours worked in the salaried sector, in full-time units, and salaried employment rate, by time relative to the first qualified healthcare worker job at a hospital. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked).

Note. Data on individuals who got their first hospital nurse job in 1988 or before are omitted from the computation.

Source. Insee, DADS panel.

Figure 3: Lifecycle profile of hospital nurses' labor supply: decompositions

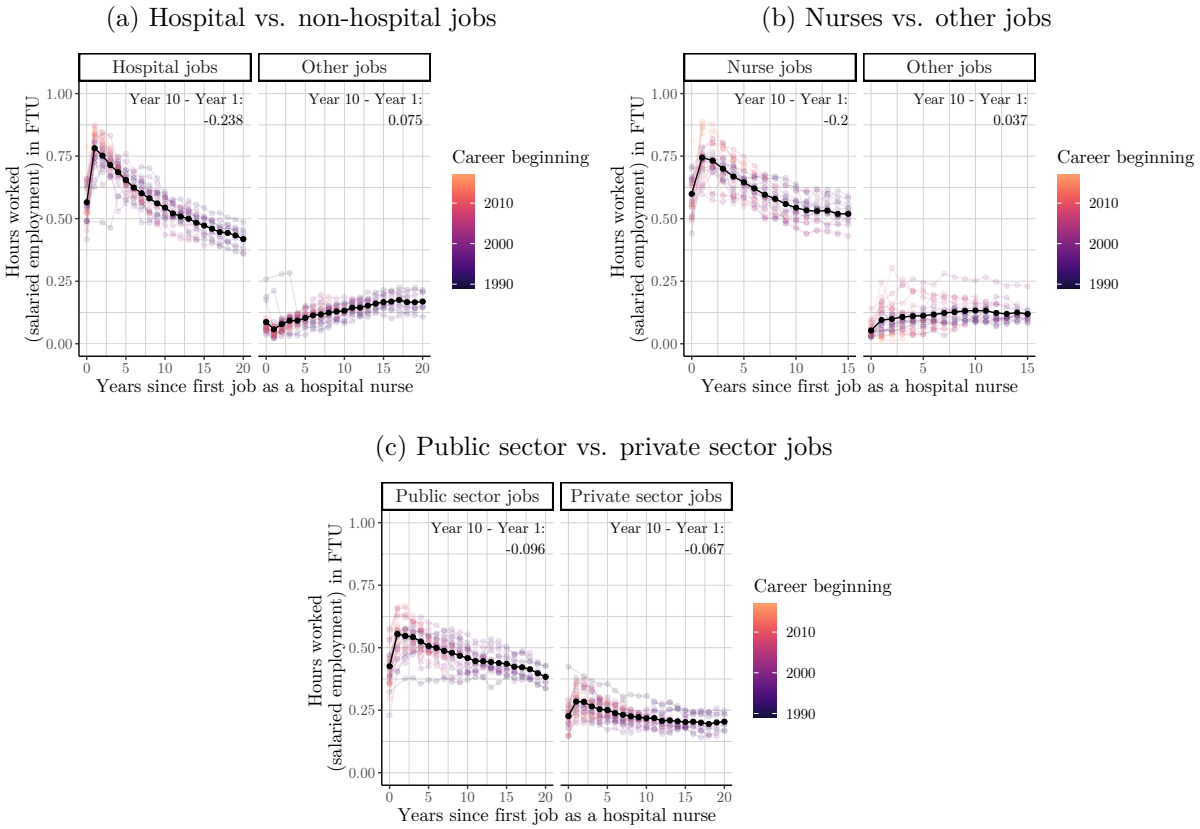

Average hours worked in the salaried sector, in full-time units, by time relative to the first qualified healthcare worker job at a hospital. Hours worked are not conditional on salaried employment in a given sector or setting, but incorporate the participation margin (0 hours worked).

Note. Data on individuals who got their first hospital nurse job in 1988 or before are omitted from the computation.

Source. Insee, DADS-EDP panel.

Figure 4: Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector

(a) Hours worked

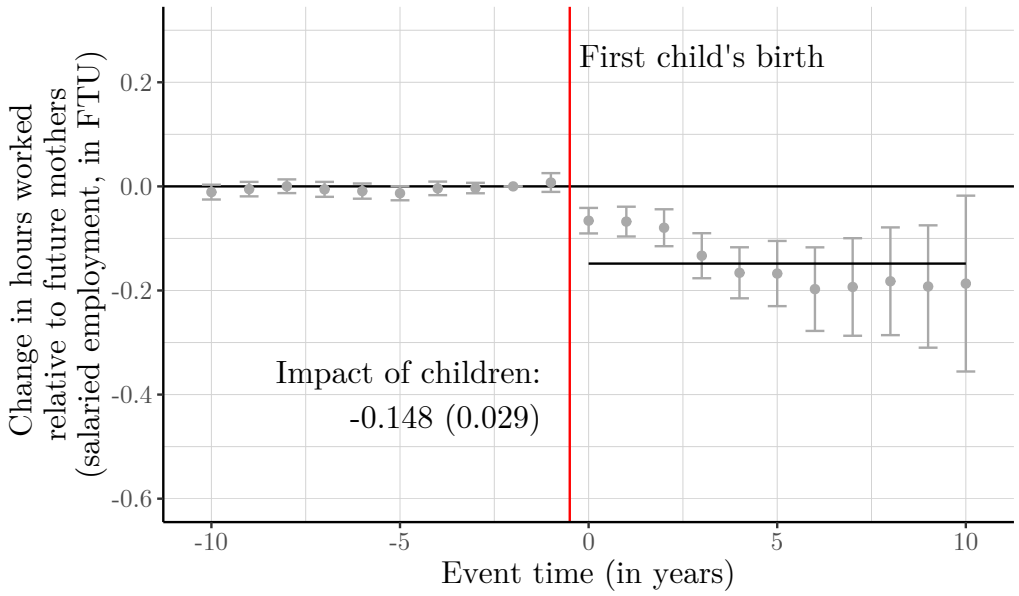

(b) Salaried employment rate

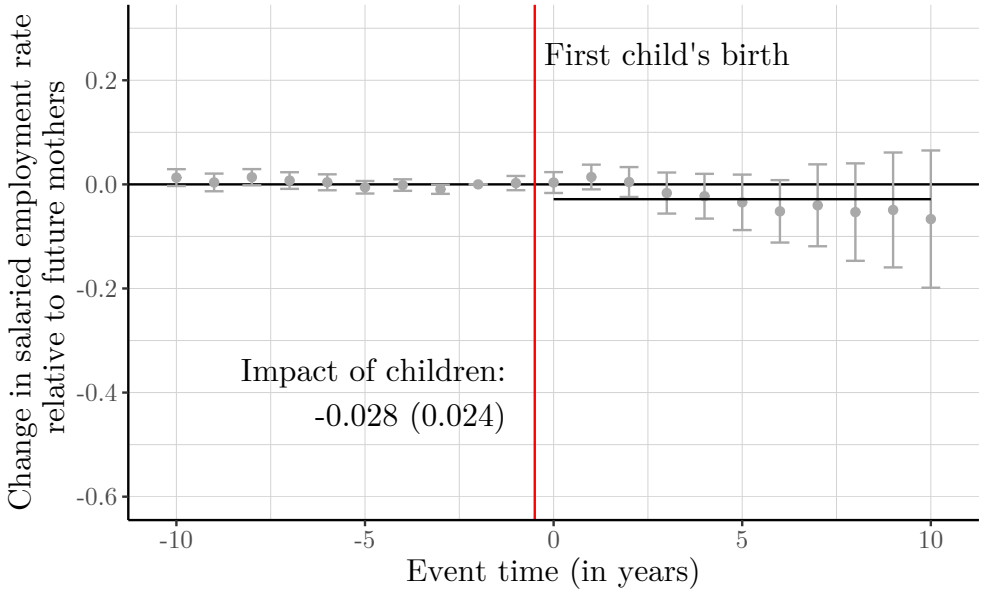

Event-study estimates of the impact of children on mothers' hours worked in the salaried sector, in full-time units, and salaried employment rate, by time since first child's birth. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked). Standard errors are clustered at the individual level and estimated by bootstrap with 200 replications.
Source. Insee, DADS-EDP panel.

Figure 5: Event-study estimates of the impact of children on mothers' labor supply: decompositions

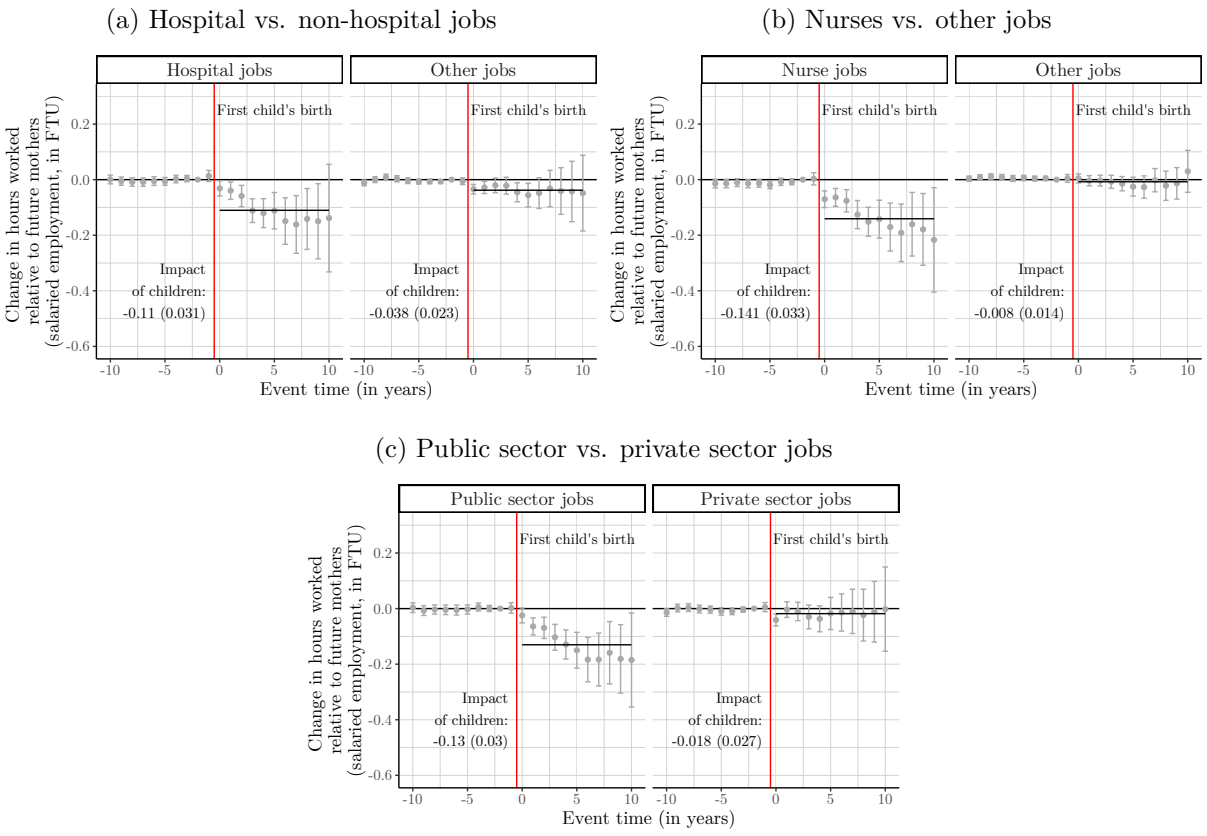

Event-study estimates of the impact of children on mothers' hours worked in the salaried sector, in full-time units, by time since first child's birth. Hours worked are not conditional on salaried employment in a given sector or setting, but incorporate the participation margin (0 hours worked). Standard errors are clustered at the individual level and estimated by bootstrap with 200 replications.
Source. Insee, DADS-EDP panel.

Figure 6: Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector

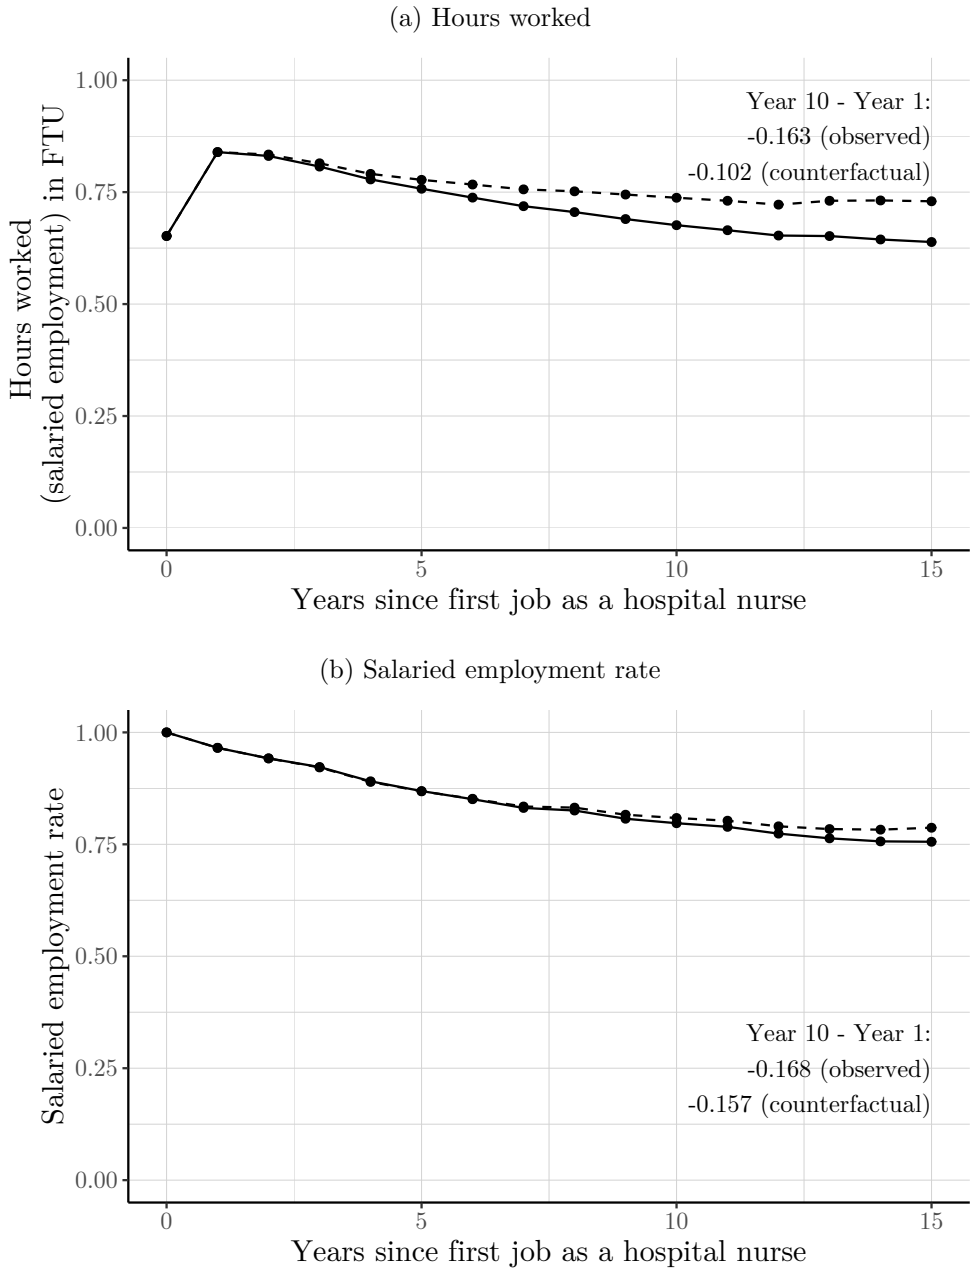

Realized and counterfactual average hours worked in the salaried sector, in full-time units, and salaried employment rate, by time relative to the first qualified healthcare worker job at a hospital. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked).
Note. Data on individuals who got their first hospital nurse job in 1988 or before are omitted from the computation.
Source. Insee, DADS-EDP panel.

Figure 7: Contribution of children to the lifecycle profile of nurses' labor supply: decompositions

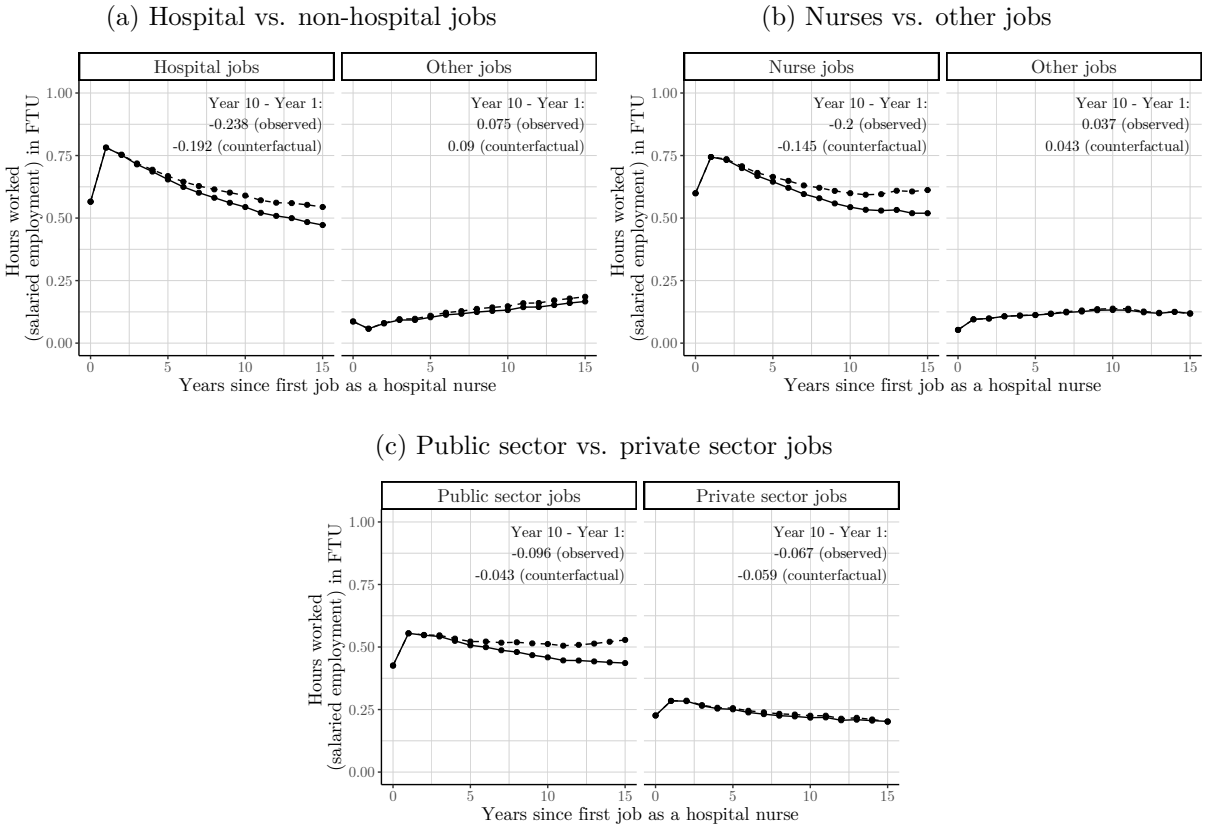

Realized and counterfactual average hours worked in the salaried sector, in FTU, by time relative to the first qualified healthcare worker job at a hospital. Hours worked are not conditional on salaried employment in a given sector or setting, but incorporate the participation margin (0 hours worked).

Note. Data on individuals who got their first hospital nurse job in 1988 or before are omitted from the computation.

Source. Insee, DADS-EDP panel.

Tables

Table 1: Detailed occupations distribution among selected jobs (2009-2017)

Detailed occupation	Share among women (in %)	Share among men (in %)
431A – Nurses managers	5.4	9.8
431B – Mental health nurses	1.9	2.9
431C – Nursery nurses	1.8	0.2
431D – Other specialized nurses	4.0	6.7
431F – General care nurses	63.8	55.6
All nurses occupations	76.9	75.2
431E – Midwives	3.5	1.0
432B – Physical therapists	1.8	3.7
432D – Other rehabilitation specialists	3.9	1.6
433A – Medical technicians	7.0	10.0
433B – Opticians and hearing aid professionals	0.0	0.0
433C – Other specialists in medical equipment	0.0	0.6
433D – Pharmacy technicians	2.6	2.4
434A – Social work managers	0.2	0.3
434B – Social work assistants	1.8	0.6
434C – Family economic counselors	0.1	0
434D – Specialized educators	0.9	2.1
434E – Instructors	0.4	1.4
434F – Specialized technical educators, workshop monitors	0.1	0.3
434G – Early childhood educators	0.3	0.1
435B – Socio-cultural and leisure animators	0.4	1.0

Source. DADS panel, Insee.

Table 2: Summary statistics

	Women	Men
# Observations	133,664	28,059
# Individuals	4,652	975
a. Age at first hospital nurse job[*]		
Mean	28.4	31.9
St.D.	8.4	9.6
b. Potential experience at first hospital nurse job[**]		
Mean	5.4	6.6
St.D.	4.5	5.1
c. Share of parents at age 45 (in %)[***]	81.5	79.4
d. Number of children[****]		
Mean	2.2	2.3
St.D.	0.9	1.1

[*] Among those who got their first hospital nurse job after 1988. [**] Among those who got their first job after 1988. Potential experience is defined as the difference between the year during which an individual get her first job as a hospital nurse, and the year during which she holds her first job whatever the industry or occupation. [***] Among those born before 1973. [****] Among those born before 1973 with at least one child. *Source.* DADS-EDP panel, Insee.

A DADS panel: labor supply measures

A.1 Hours worked: concept

In the DADS dataset, hours worked refers to hours for which the worker is paid under their labor contract. The data on hours is reported by employers when they fill out payroll tax forms. Before making the data available, Insee performs three checks:

- the total number of hours for a given individual × employer × year observation should not exceed an industry-specific threshold of 2,500 hours per year in a small subset of industries (mostly manufacturing industries, transportation, hotels and restaurants), and 2,200 hours per year elsewhere;
- the implied hourly wages should exceed 80% of the minimum wage;
- the total number of hours should be positive, with the exception of a narrow subset of occupations (mostly journalists and salespersons) working on a fixed-price or commission basis.

If one of these conditions is not met, Insee ascribes hours to the observation to make the hourly wage consistent within narrow cells defined by 4-digit occupation, full-time or part-time status, age and gender.

During a maternity leave, as an employee is not paid by for any hours by her employer but is instead paid by the Social Security (and may receive a top-up payment from her employer), hours worked are equal to 0. Workers not paid by the hour are an exception to this rule because their hours are imputed based on days paid, which do not vary during maternity leave. As a result, the DADS dataset overestimates hours paid – and underestimates hourly wages – for such workers during years when they give birth to children. In general, these workers belong to the "Manager and professionals" occupation group, so that this is not a concern for this particular paper.

A.2 Full-time units conversion

Hours worked are converted in full-time units using a very simple approach, that relies on three variables: working-time status, days worked and lastly hours worked. This approach caps time worked at 1 for individuals who work full-time during an entire year, so that it does not incorporate overtime. The main advantage of this method is that it allows to compare time worked even when the legal duration of work changes, which is the case over my time period of interest as this duration changed from 39h to 35h per week in the beginning of the 2000s.

Full-time workers who are observed to be employed an entire year are assigned 1 full-time unit for this year. Full-time workers who are not observed to be employed for an entire year are assigned a value that is proportional to days worked, so that it would be 1 if they were working for the entire year.[8] Part-time workers are assigned a value proportional to hours worked, so

[8] In the DADS data, an entire year corresponds to 360 days worked.

that their time worked would be 1 if their hours worked matched the legal duration of work over an entire year.

An issue with this approach is that hours worked are not observed before 1995 in the DADS data. Before this date, only the working-time status and days worked can be observed. I choose to impute full-time units to part-time workers under the assumption that they are on a 50% schedule, so that their time worked is proportional to days worked and equal to 0.5 if they work an entire year. I make this choice because 50% schedule was likely the most frequent choice among part-time workers. Figure A.1 makes it salient by plotting the distribution of hours worked, relative to (i) the legal duration of work and (ii) days worked, among part-time job spells observed between 1995 and 1998 for workers who at one point in time were observed in hospital nurses positions.

Figure A.1: Distribution of hours worked for part-time workers, 1995-1998

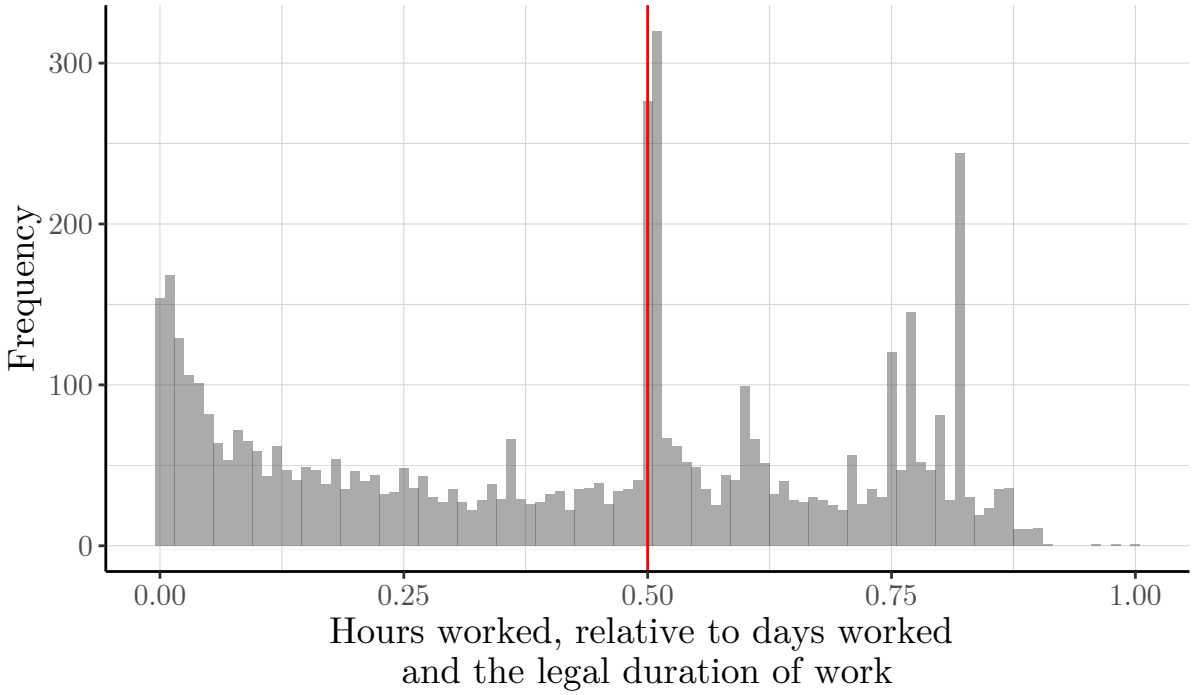

Distribution of hours worked divided by the legal annual duration of work for full-time workers, among individuals who hold a hospital nurse job at least once.
Source. Insee, DADS panel.

B How well does the sample cover the nurse occupation?

Table 1 shows that even though I cannot rely on the detailed occupation variable to delineate my sample, it mostly covers individuals that can reasonably be considered as hospital nurses. Reciprocally, a question is how many nurses do, at some point of their lives, hold a job as a hospital nurse, and thus fall within the universe that I intend to cover. To investigate this issue, I consider all individuals (i) observed as holding a salaried nurse job, in terms of the detailed occupation variable, at any point in time since 2010; and (ii) whose first job, whatever the industry or occupation, started in 2010 or later. The second condition is meant to capture individuals whose entire career can be observed in terms of the detailed occupation variable, so as to avoid selection issue that would stem from nurses gradually leaving the occupation over time (see Figure 3). Because nurses cannot begin their careers as freelance nurses, but have to hold a job as a salaried nurses, these conditions should include all individuals who began their careers in 2010 or after and hold a nurse job at some point. I then compute the share of these individuals who fall within my sample. This measure gives a lower bound of the share of nurses who hold a hospital nurse job at some point of their lives, given that I can only observe the very beginning of a career for the selected individuals.

Figure B.1 displays my estimates. Even though there is variation from one cohort to the other, which is probably due to small sample size, it suggests that at least 74% of nurses hold a job as a nurse at a hospital at some point of their lives.

Figure B.1: Share of nurses who ever hold a job at a hospital

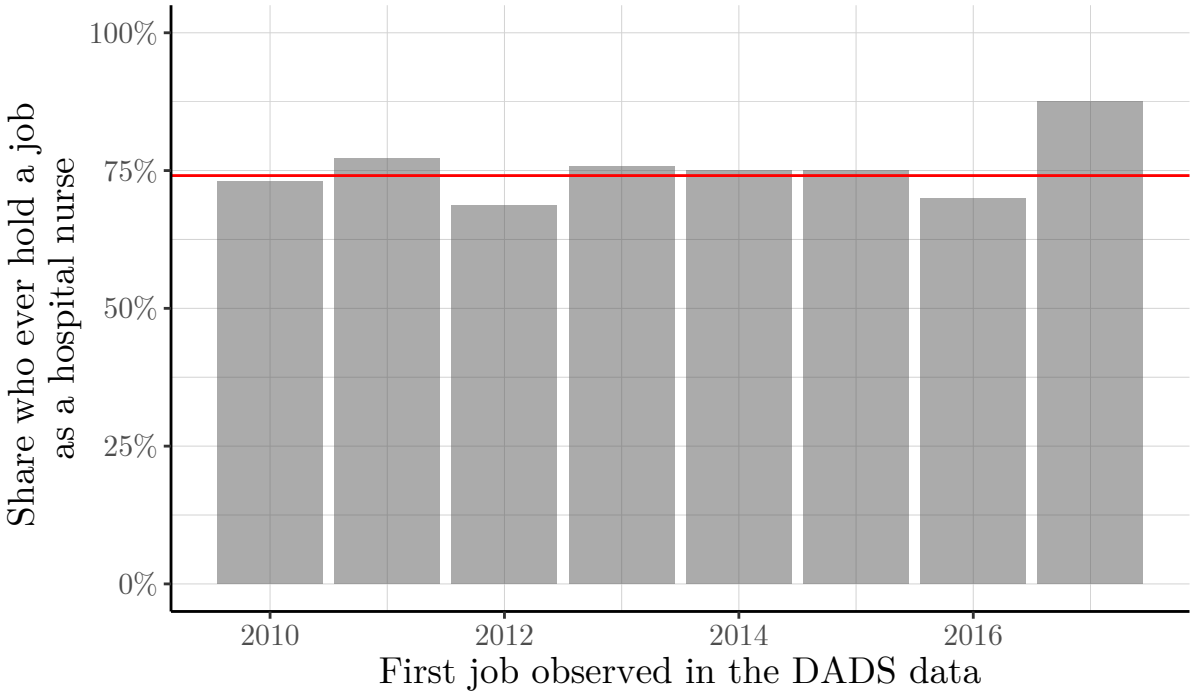

Share of nurses who began their careers in 2010 or later who are observed to hold a job as a hospital nurse at some point of their lives.
Source. Insee, DADS panel.

C Additional results

C.1 Lifecycle profiles

Figure C.1: Lifecycle profile of hospital nurses' labor supply: decompositions at the extensive margin

(a) Hospital and nurse jobs jobs

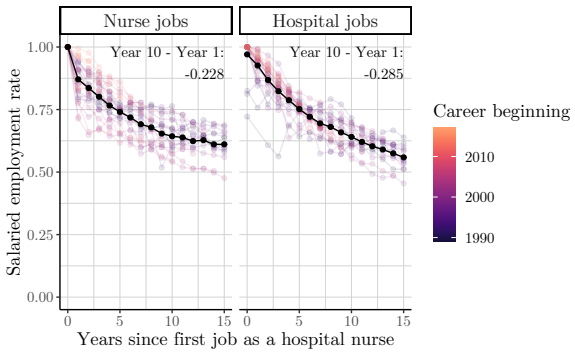

(b) Public sector vs. private sector jobs

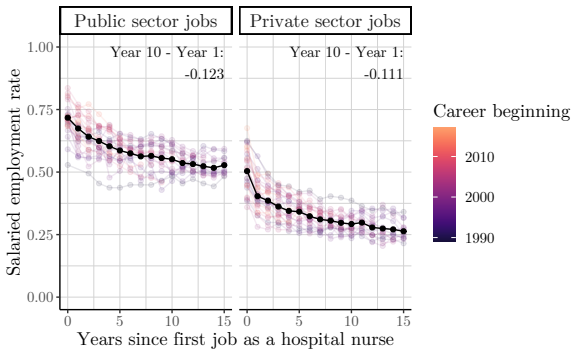

Salaried employment rate, by time relative to the first qualified healthcare worker job at a hospital.
Note. Data on individuals who got their first hospital nurse job in 1988 or before are omitted from the computation.
Source. Insee, DADS-EDP panel.

C.2 Children-related labor supply decisions

C.2.1 Extensive and intensive margins of fertility decisions

My approach that is based on the child penalty framework of Kleven, Landais, and Søgaard (2019) identifies the causal effect of motherhood on mother's labor outcomes. However, this causal effect mixes up the consequences of two interventions: (i) becoming a mother, as opposed to remaining childless, which correspond to the extensive margin of fertility decisions; and (ii) having additional children for women who are with child, which is the intensive margin of fertility decision. Because they are not conditional on subsequent fertility decisions, my estimates of the impact of motherhood will therefore incorporate the causal effect of both margins. Specifically, the short-run effect of motherhood are likely to reflect the short-run consequences of the extensive margin, because it is uncommon to have additional children the very year one's firt child is born. By contrast, my long-run estimates will mix up (i) the long-run consequences of the extensive margin and (ii) the short-run effect of the intensive margin, with weights that depend on the timing and frequency of subsequent childbirths. This is especially true in this context, in which most mothers choose to have more than one child.

To gain further insights into this issue, I replicate my event-study analysis (a) only relying on observations related to mothers of one or mothers of additional children *at least two years before their second child is born*; and (b) considering the second child's birth instead of the first child's birth as my event of interest *and* only relying on observations related to mothers of two or mothers of additional children *at least two years before their third child is born*. Under parallel trends and limited anticipation assumptions very similar to Assumptions 1 and 2,[9] this allows to distinguish between (a) the dynamic effect of the first child and (b) the dynamic effect of having one additional child (i.e. the second child), i.e. the dynamic effect of (a) the extensive and (b) the intensive margin of fertility, without contamination from subsequent fertility decisions.

Figure C.2 displays these estimates. While they may be slightly less precise than the baseline estimates, because they rely on less observations, they are still informative to some extent. Specifically, they show that while both margins seem to have somewhat similar short-run effects, the impact of having one additional child is quite short-lived. In other words, the impact of the second child seems to vanish after a few years, whereas the impact of the first one is long-lasting. Moreover, the magnitude of the impact of the first child is not very different from that of my baseline estimates that mix up all margins.

This suggests that these baseline results are mostly driven by the long-run consequences of fertility decisions at the extensive margin, i.e. the decisions to become a mother, rather than the short-run consequences of fertility decisions at the intensive margin, i.e. the decision to have one additional child among women who are already with child. These results are also, to some extent, informative about the origins of such labor supply decisions. Indeed, if these decisions were purely the result of children-related time constraint, as opposed to norms and preferences, one would expect the dynamic path of the effect of the first child to be very similar

[9] Specifically, I assume that (i) mothers that are to have a second child are a good comparison group for those who just had their second child and (ii) the second and the third child have no effect on their mother's labor supply up until one year before they are born.

to that of the second child, because conditional on their age, the needs of children should not be strongly dependent on their rank among their siblings. That they do really differ therefore seems to indicate either (i) large returns to scale in the children production function or (ii) that motherhood-related time allocation decisions involve gender norms and preferences.

Figure C.2: Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector, by fertility margin

(a) Extensive margin (first child)

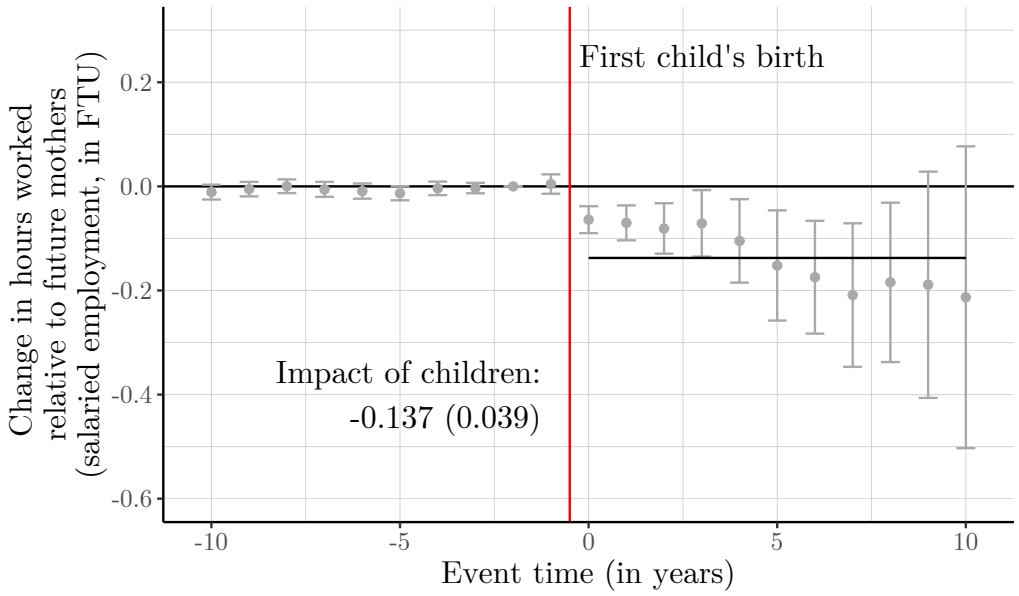

(b) Intensive margin (second child)

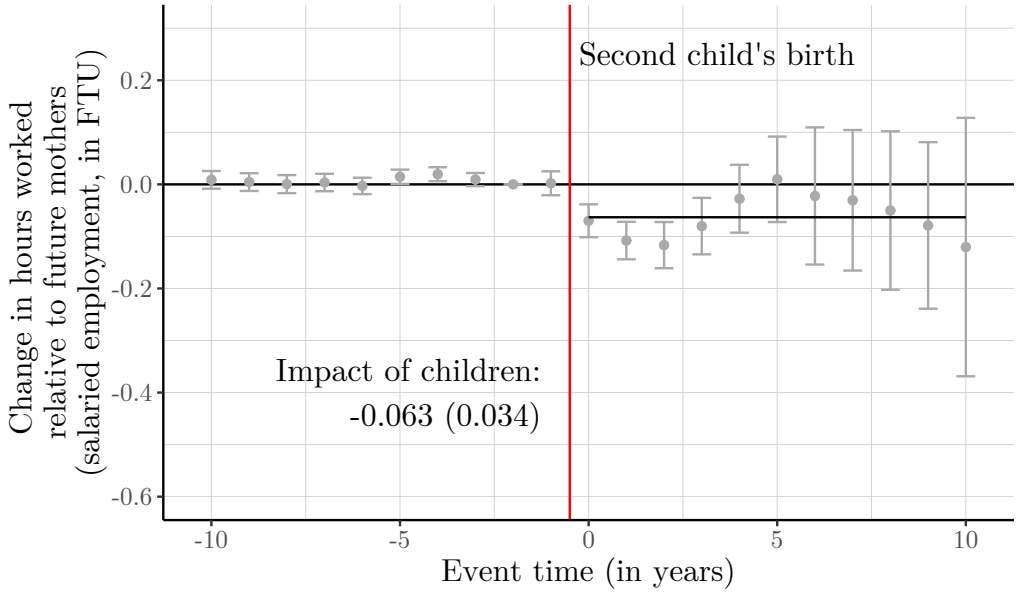

Event-study estimates of the impact of children on mothers' hours worked in the salaried sector, in full-time units, by time since first child's birth. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked). Standard errors are clustered at the individual level and estimated by bootstrap with 200 replications.

Source. Insee, DADS-EDP panel.

C.2.2 Other margins of labor supply

Figure C.3: Event-study estimates of the impact of children on mothers' labor supply: decompositions at the extensive margin

(a) Hospital and nurse jobs

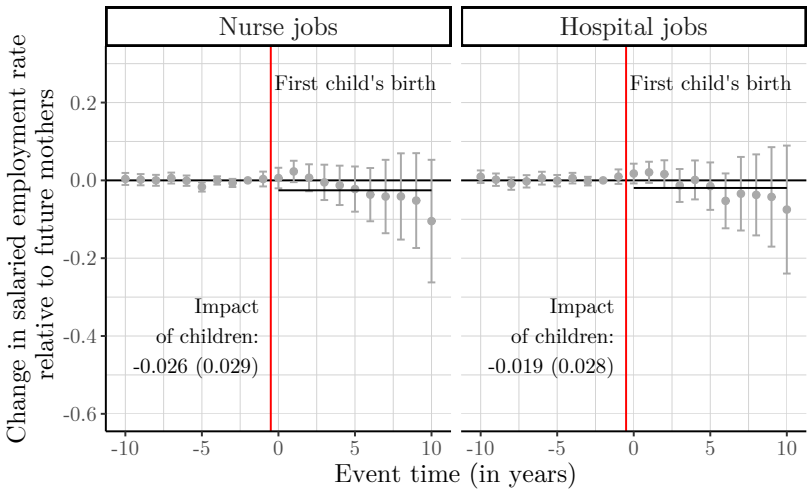

(b) Public sector vs. private sector jobs

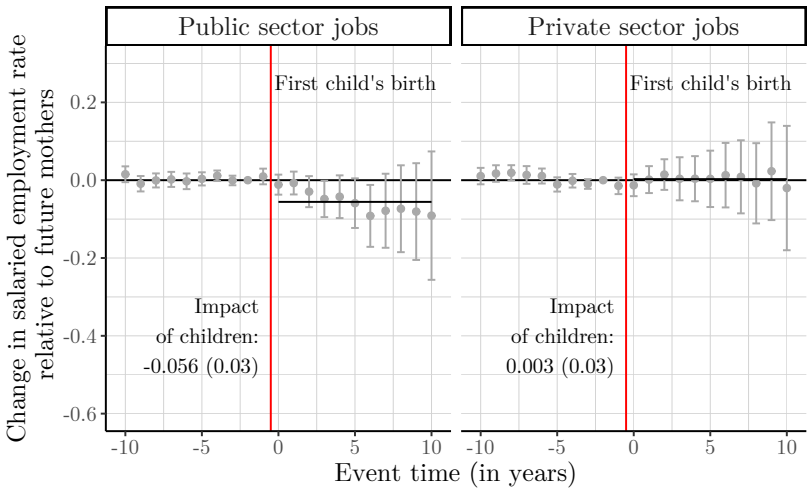

Event-study estimates of the impact of children on mothers' participation in the salaried sector, by time since first child's birth. Standard errors are clustered at the individual level and estimated by bootstrap with 200 replications.
Source. Insee, DADS-EDP panel.

Figure C.4: Event-study estimates of the impact of children on mothers' labor outcomes: labor earnings

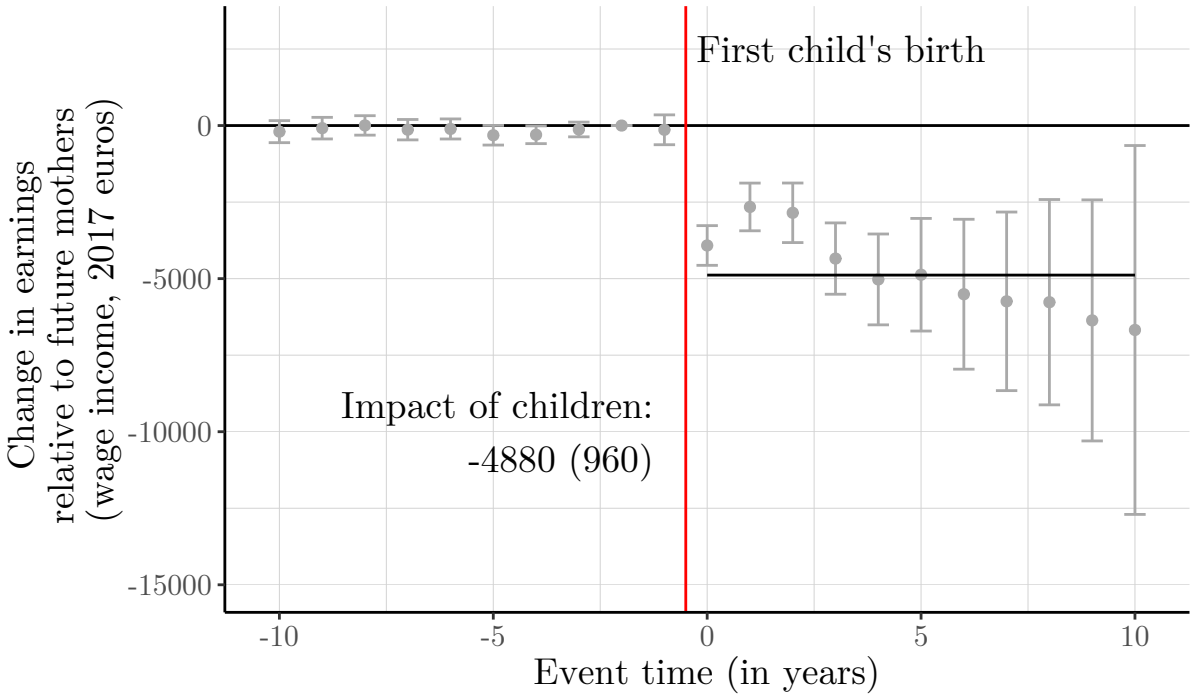

Event-study estimates of the impact of children on mothers' salaried labor earnings, by time since first child's birth. Standard errors are clustered at the individual level and estimated by bootstrap with 200 replications.

Source. Insee, DADS-EDP panel.

C.2.3 Children-related labor supply decisions: fathers

Figure C.5: Event-study estimates of the impact of children on fathers' labor supply: total labor supply in the salaried sector

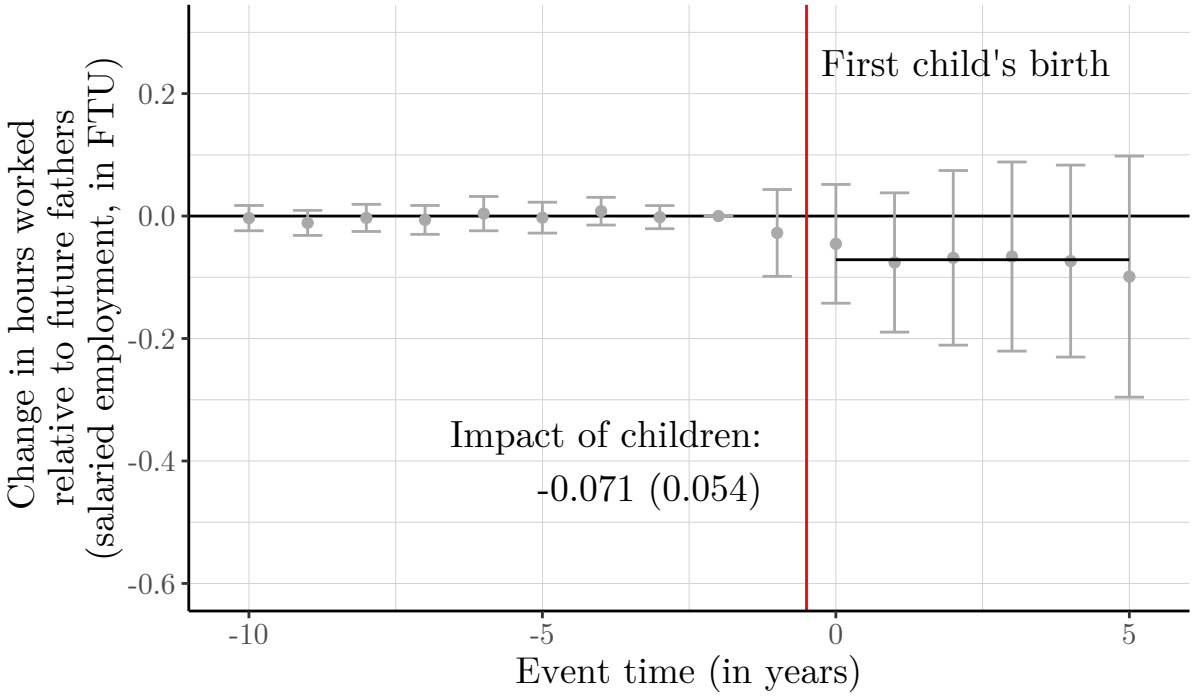

Event-study estimates of the impact of children on fathers' hours worked in the salaried sector, in full-time units, by time since first child's birth. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked). Standard errors are clustered at the individual level and estimated by bootstrap with 200 replications.
Source. Insee, DADS-EDP panel.

C.3 Overtime, working conditions and hourly wages

Hours worked, as measured in full-time units, does not include overtime because this measure is capped to 1 for full-time workers. As such, if some nurses choose to keep their full-time jobs, but cut their overtime hours upon becoming mothers, then my estimates of the magnitude of children-related labor supply decisions will underestimate how important these adjustments are. Additionally, nursing is characterized by particularly salient constraints regarding working time: because healthcare has to be provided continuously, shift work is a common working time arrangement, which is not always sufficient to prevent unforeseeable planning changes. As a result, nurses may be prone to turn to jobs that are less exposed to these time constraints, and offer better work-family conciliation upon becoming mothers. This margin would not appear in hours worked changes measured in full-time units, as it involves other dimensions of working time.

To investigate these dimensions, in Figure C.6 I replicate my event-study analysis, focusing on two new outcomes. First, I consider hours worked including overtime, relative to an entire year of full-time work.[10] This measure is akin to hours worked measured in full-time units, except that it is no longer capped to 1 for full-time workers. My results are extremely similar to those obtained with hours worked measured in full-time units, which indicates that changes in overtime are not an important margin for children-related labor supply decisions.

Secondly, I consider hourly wages. The rationale for this choice is that pay is extremely rigid, especially in the public sector where the baseline wage rate is almost uniformly set as a function of tenure. As a result, conditional on tenure individual differences in hourly wages are almost entirely driven by difference in (i) hours worked, as overtime hours are paid higher than other hours; and (ii) various premiums and bonuses that are tightly linked to the work setting, e.g. shift work or night work. As such, children-related changes in motherhood would be indicative of changes in working conditions. I find that the effect of children on the wage rate is a quite precisely estimated 0 (the baseline wage rate being around €15), except for the very short run. As a result, the compensating differentials that such adjustments would involve do not seem to be at play.

[10] I restrict myself to the 1995-2017 time-period for during which hours are actually observed in the DADS data.

Figure C.6: Event-study estimates of the impact of children on mothers' labor outcomes: hours worked including overtime and hourly wages

(a) Hours worked, including overtime

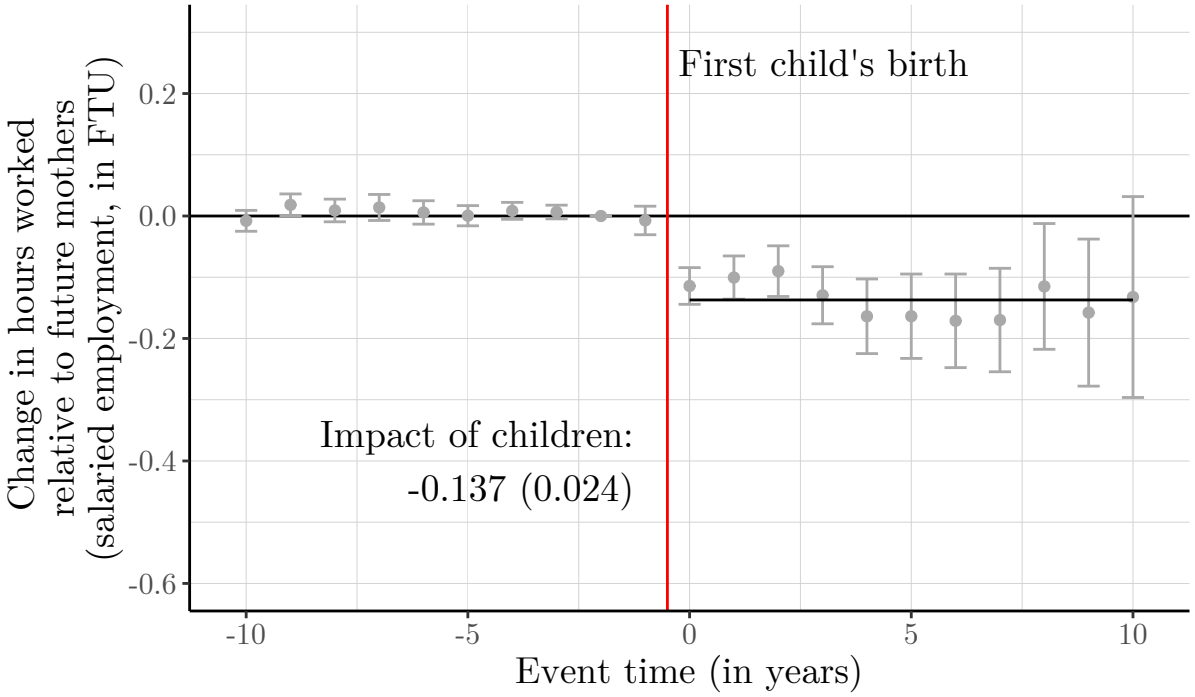

(b) Hourly wages

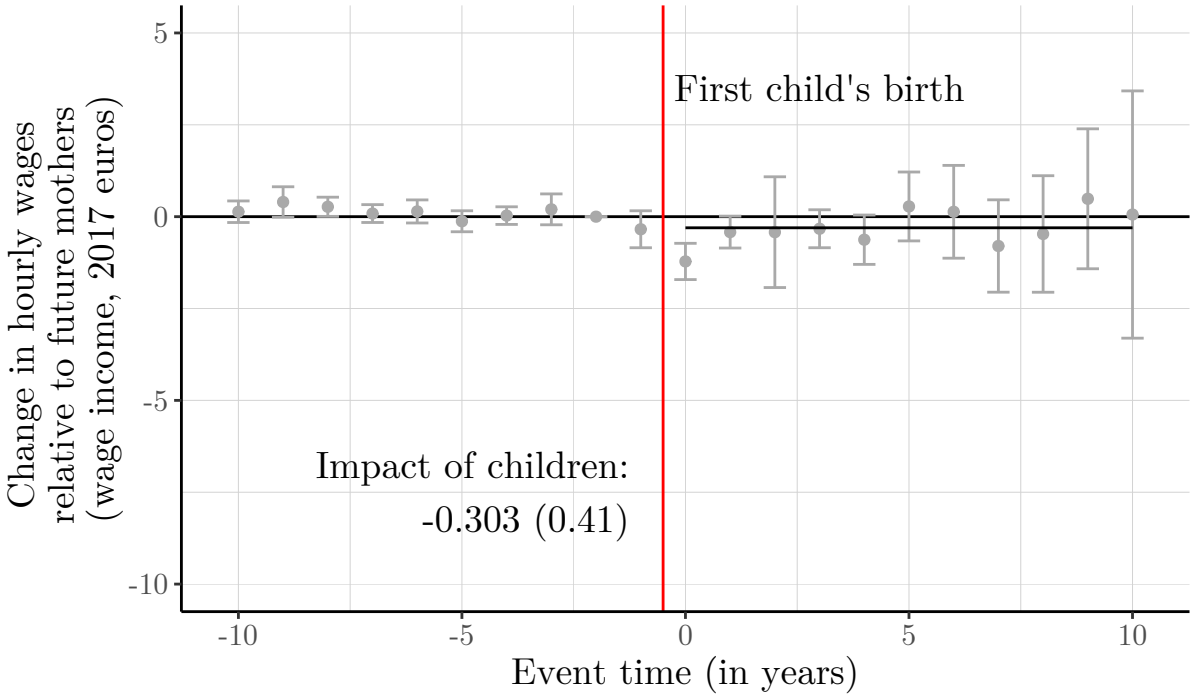

Event-study estimates of the impact of children on mothers' salaried labor outcomes, by time since first child's birth. Hours worked include overtime and are measured relative to the median number of hours for full-time workers who work and entire year (2028 hours per year until 2001, 1820 hours since 2002). Standard errors are clustered at the individual level and estimated by bootstrap with 200 replications.
Source. Insee, DADS-EDP panel.

C.4 Contribution to lifecycle profiles of labor supply

Figure C.7: Contribution of children to the lifecycle profile of nurses' labor supply: decompositions at the extensive margin

(a) Hospital and nurse jobs jobs

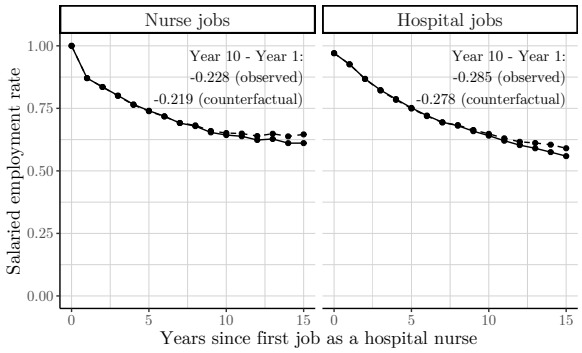

(b) Public sector vs. private sector jobs

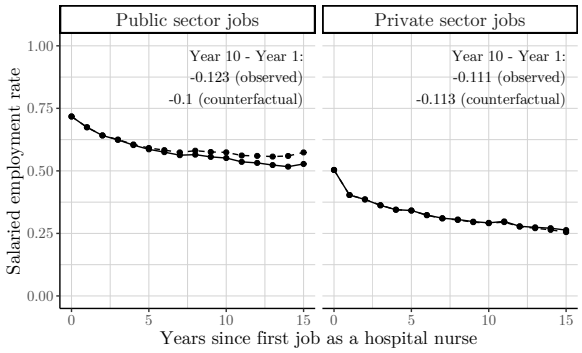

Realized and counterfactual salaried employment rate, by time relative to the first qualified healthcare worker job at a hospital.

Note. Data on individuals who got their first hospital nurse job in 1988 or before are omitted from the computation.

Source. Insee, DADS-EDP panel.

D Robustness checks

D.1 Left-censoring issue

Figure D.1: Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector

(a) Hours worked

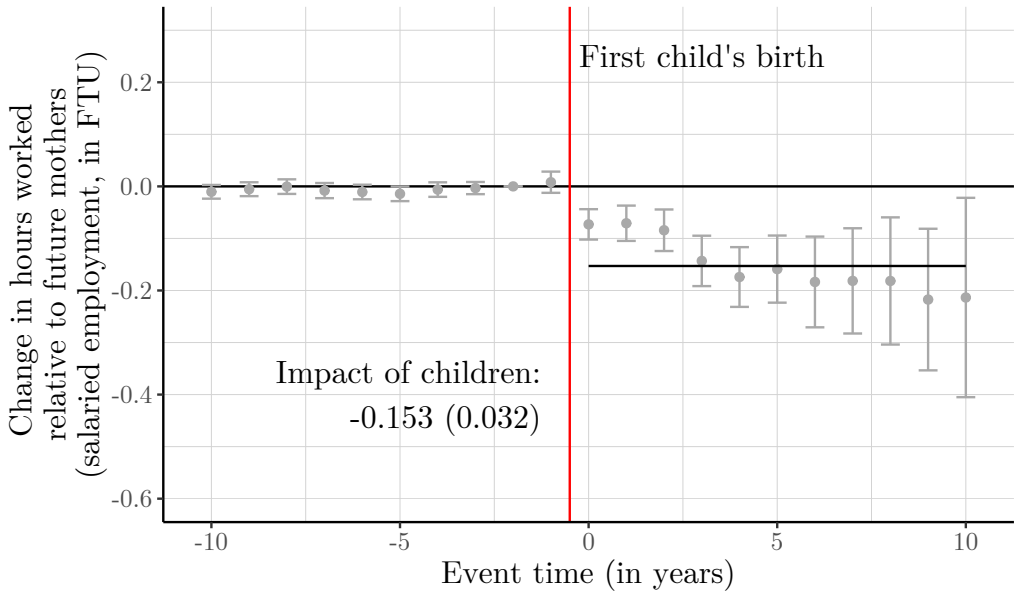

(b) Salaried employment rate

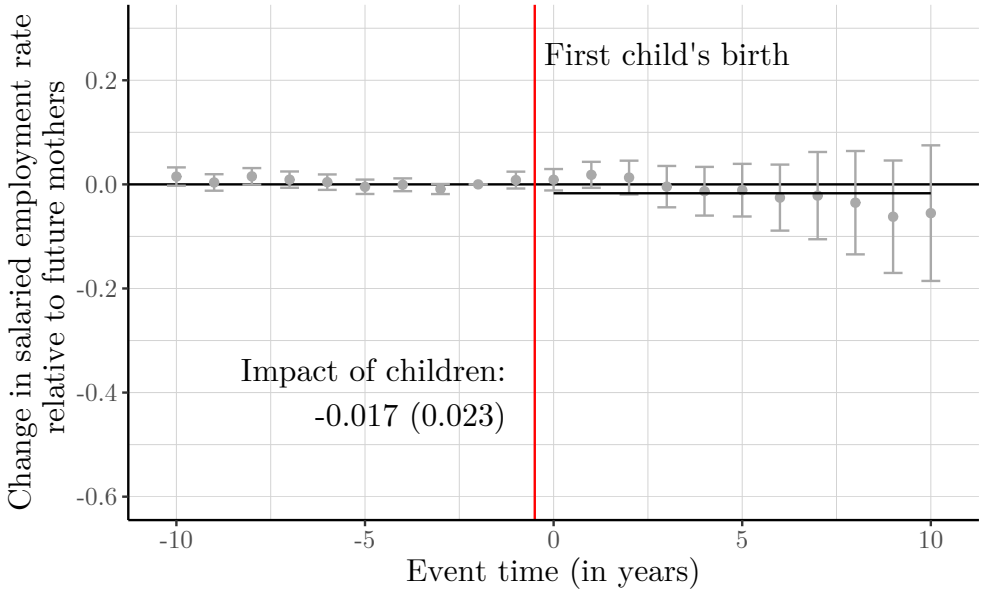

Event-study estimates of the impact of children on mothers' hours worked in the salaried sector, in full-time units, and salaried employment rate, by time since first child's birth. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked). Standard errors are clustered at the individual level and estimated by bootstrap with 200 replications.
Source. Insee, DADS-EDP panel.

Figure D.2: Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector

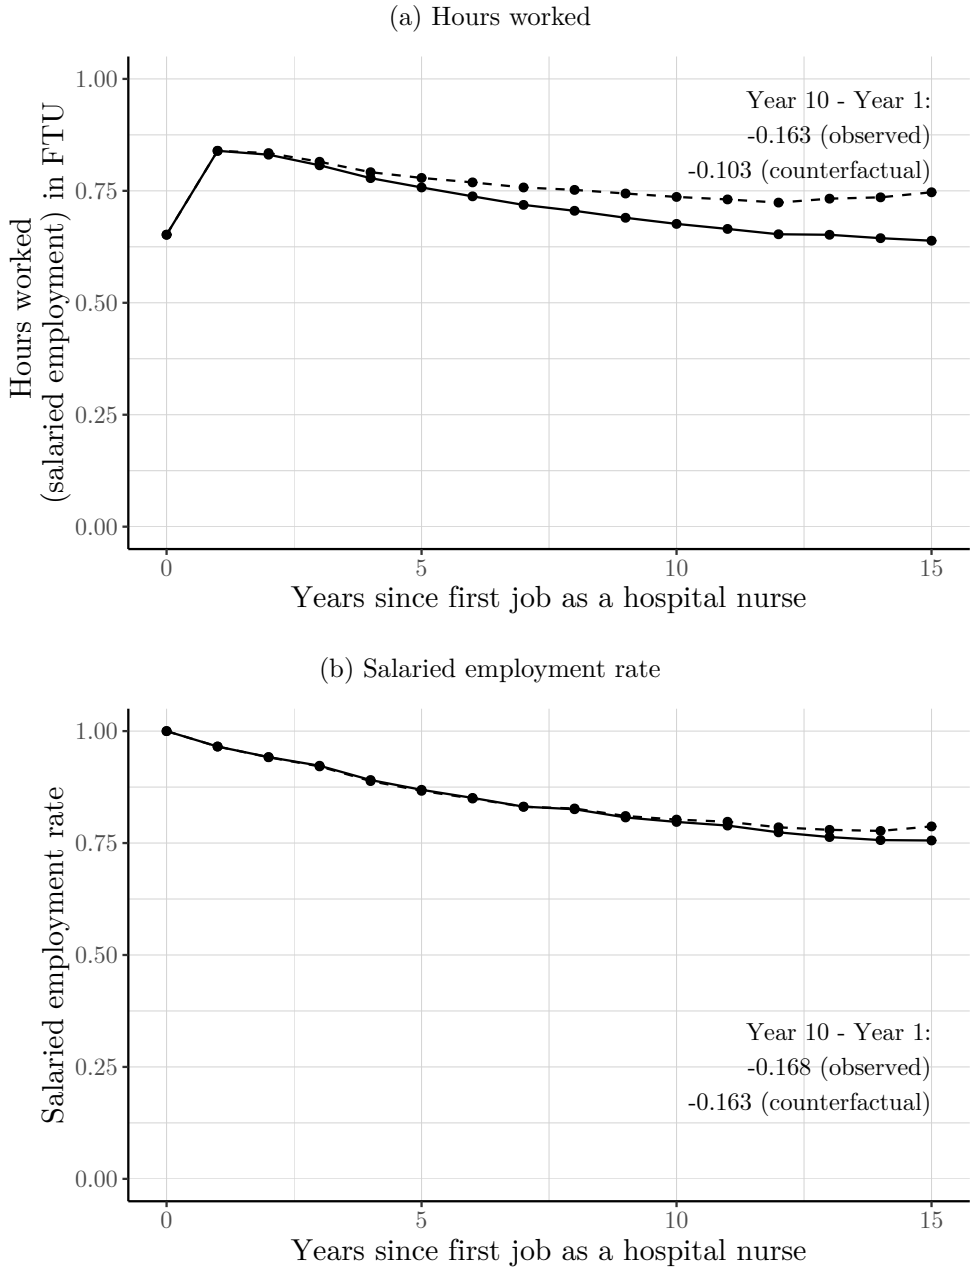

Realized and counterfactual average hours worked in the salaried sector, in full-time units, and salaried employment rate, by time relative to the first qualified healthcare worker job at a hospital. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked).
Note. Data on individuals who got their first hospital nurse job in 1988 or before are omitted from the computation.
Source. Insee, DADS-EDP panel.

D.2 Right-censoring issue

Figure D.3: Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector

(a) Hours worked

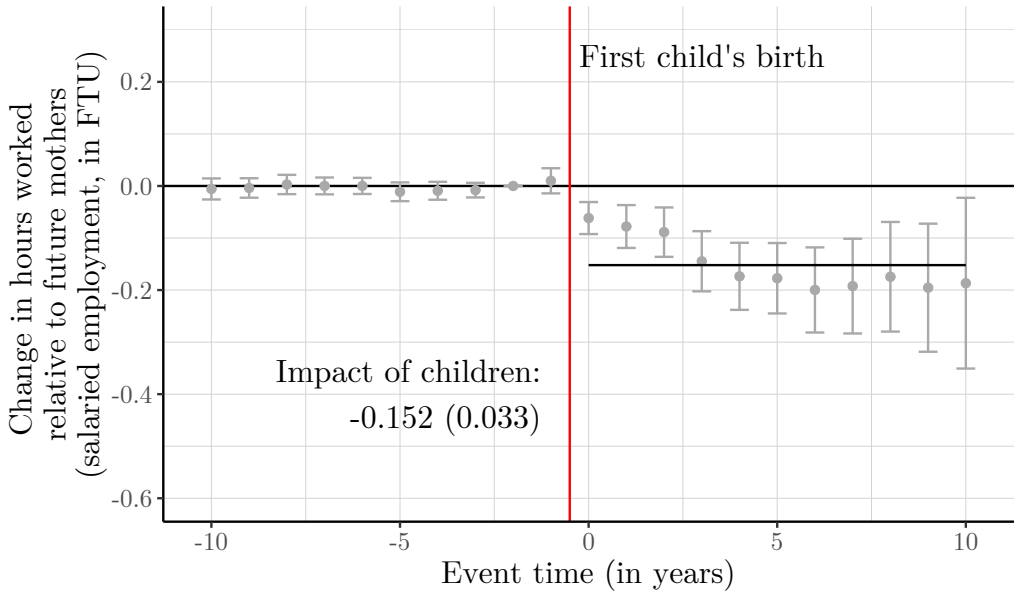

(b) Salaried employment rate

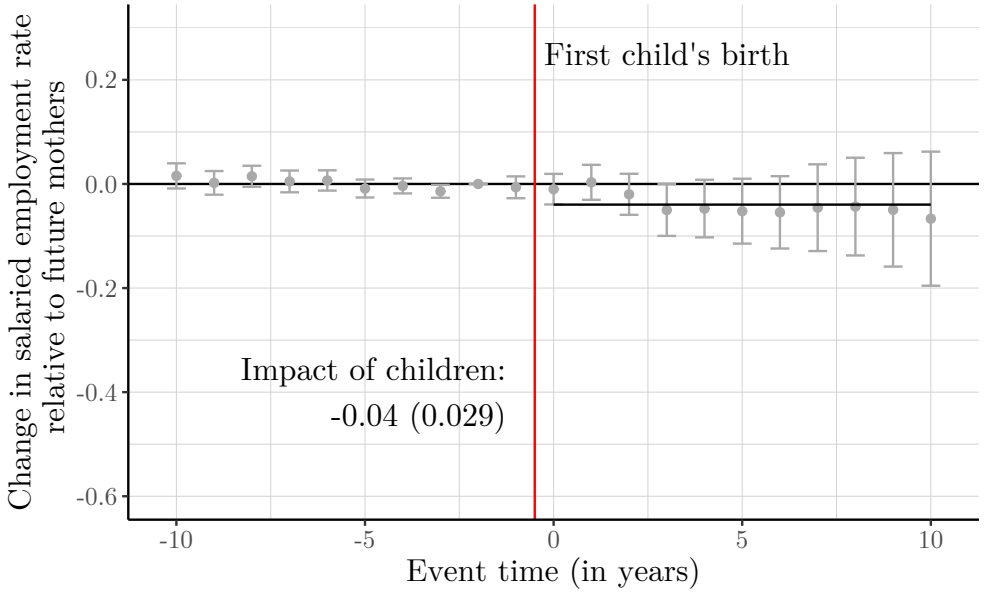

Event-study estimates of the impact of children on mothers' hours worked in the salaried sector, in full-time units, and salaried employment rate, by time since first child's birth. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked). Standard errors are clustered at the individual level and estimated by bootstrap with 200 replications.
Source. Insee, DADS-EDP panel.

Figure D.4: Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector

(a) Hours worked

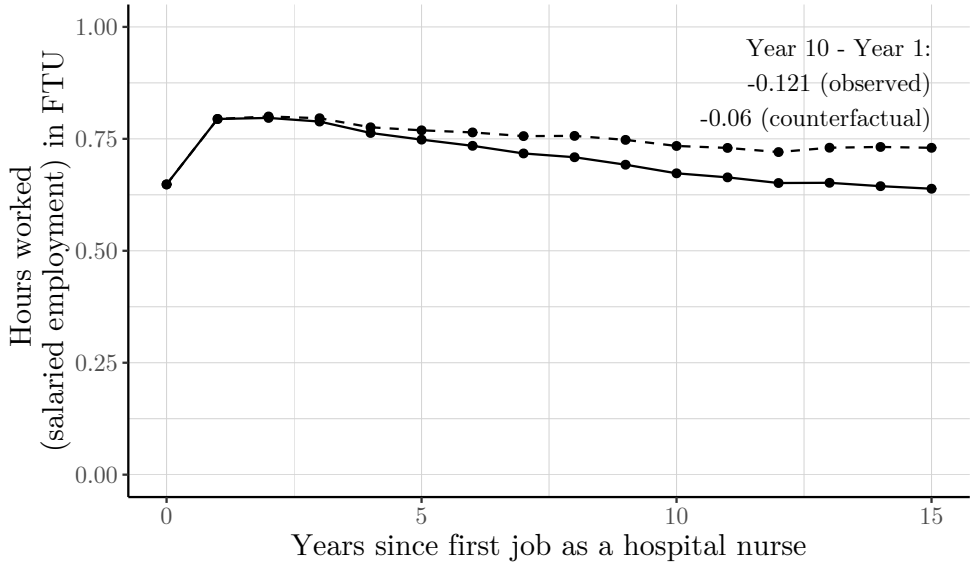

(b) Salaried employment rate

Realized and counterfactual average hours worked in the salaried sector, in full-time units, and salaried employment rate, by time relative to the first qualified healthcare worker job at a hospital. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked).
Note. Data on individuals who got their first hospital nurse job in 1988 or before are omitted from the computation.
Source. Insee, DADS-EDP panel.

D.3 Hours worked measurement

Figure D.5: Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector

(a) Hours worked

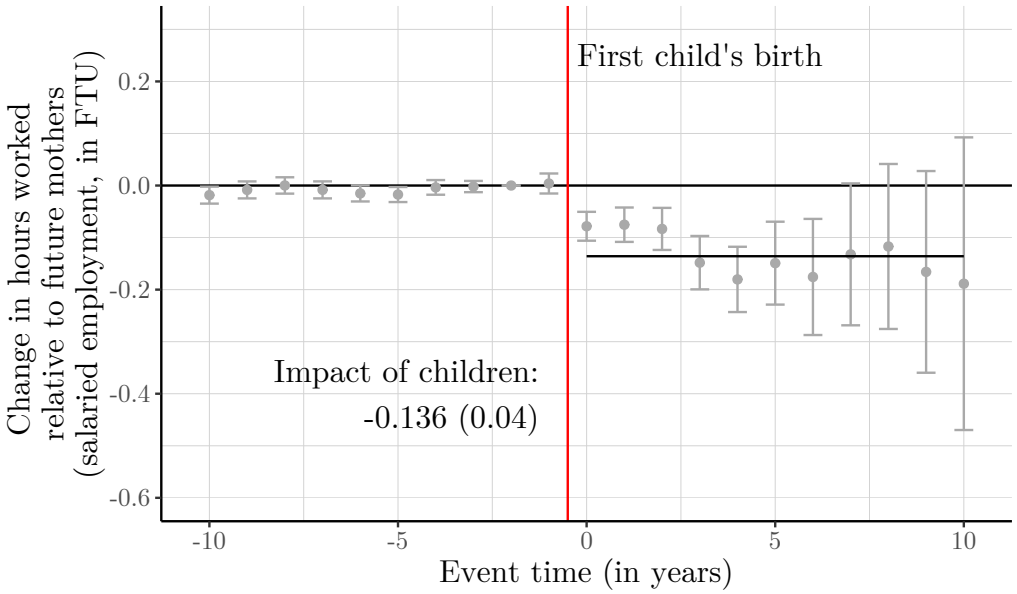

(b) Salaried employment rate

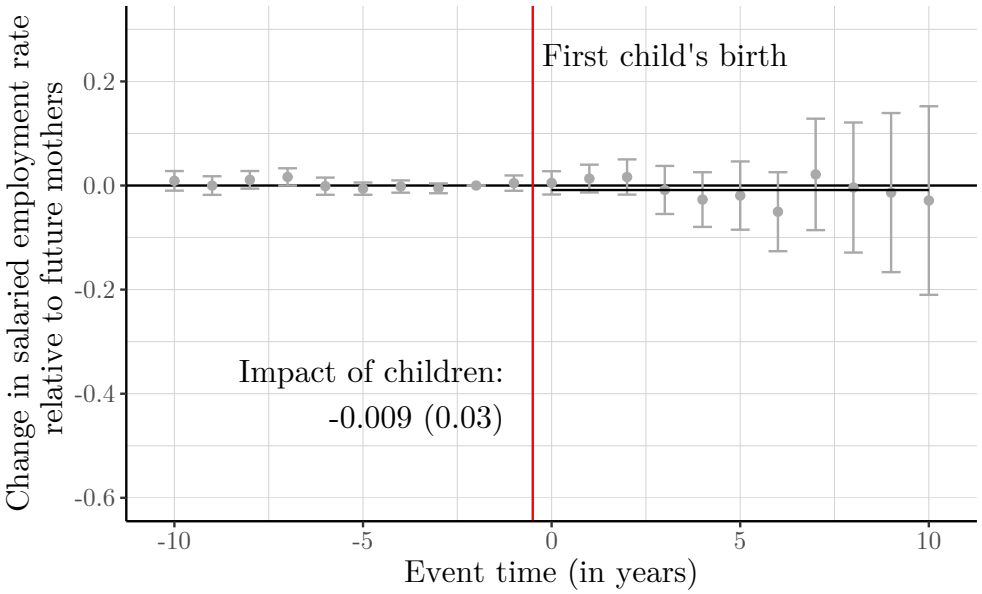

Event-study estimates of the impact of children on mothers' hours worked in the salaried sector, in full-time units, and salaried employment rate, by time since first child's birth. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked). Standard errors are clustered at the individual level and estimated by bootstrap with 200 replications.
Source. Insee, DADS-EDP panel.

Figure D.6: Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector

(a) Hours worked

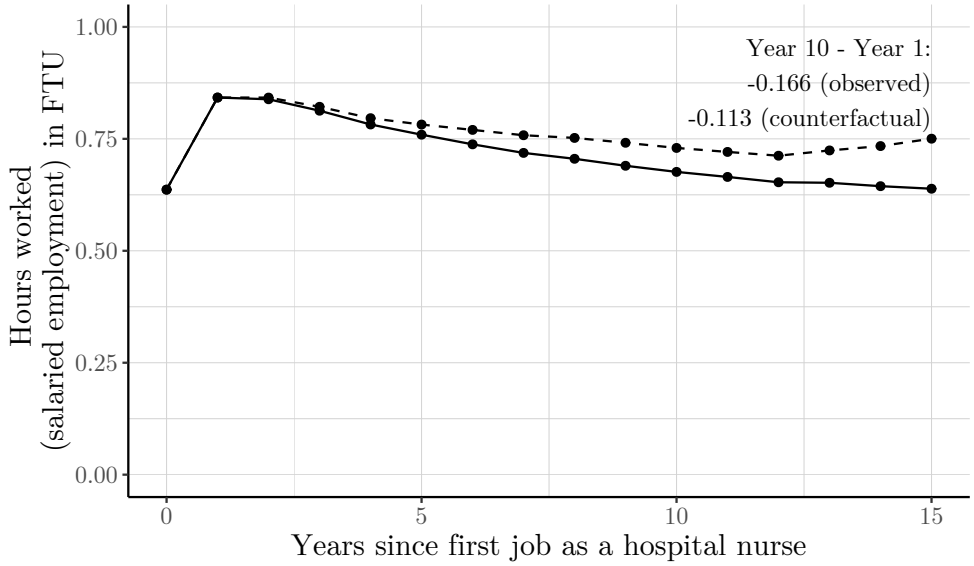

(b) Salaried employment rate

Realized and counterfactual average hours worked in the salaried sector and, in full-time units, and salaried employment rate by time relative to the first qualified healthcare worker job at a hospital. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked).

Note. Data on individuals who got their first hospital nurse job in 1988 or before are omitted from the computation.

Source. Insee, DADS-EDP panel.

D.4 Stable control group

Figure D.7: Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector

(a) Hours worked

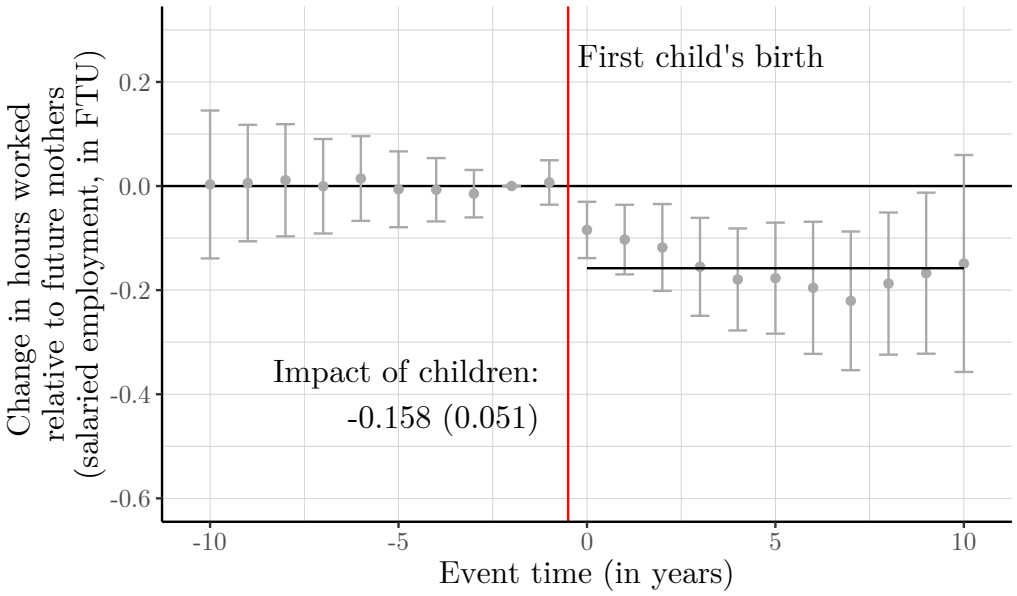

(b) Salaried employment rate

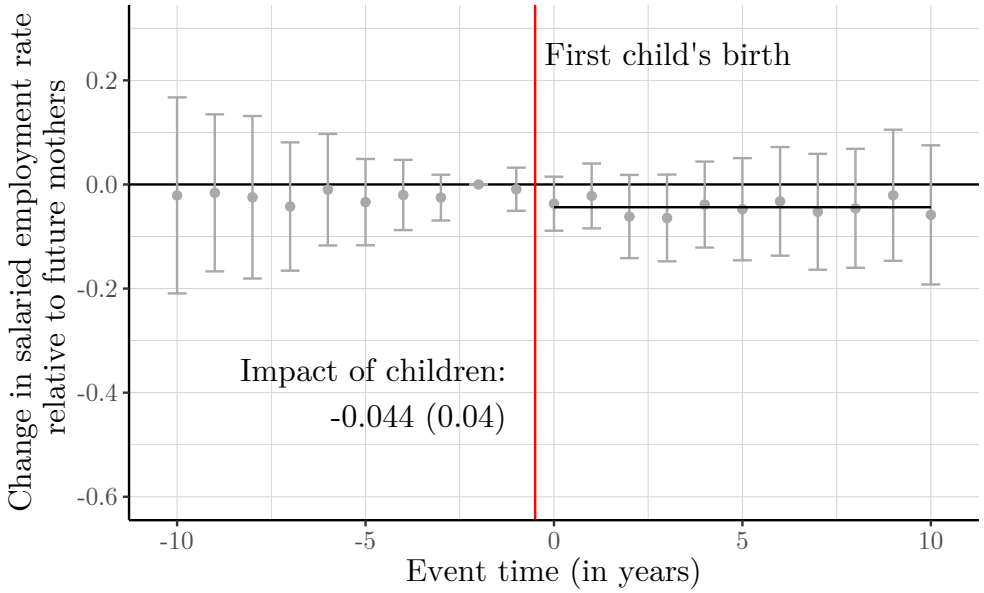

Event-study estimates of the impact of children on mothers' hours worked in the salaried sector, in full-time units, and salaried employment rate, by time since first child's birth. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked). Standard errors are clustered at the individual level and estimated by bootstrap with 200 replications.
Source. Insee, DADS-EDP panel.

Figure D.8: Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector

(a) Hours worked

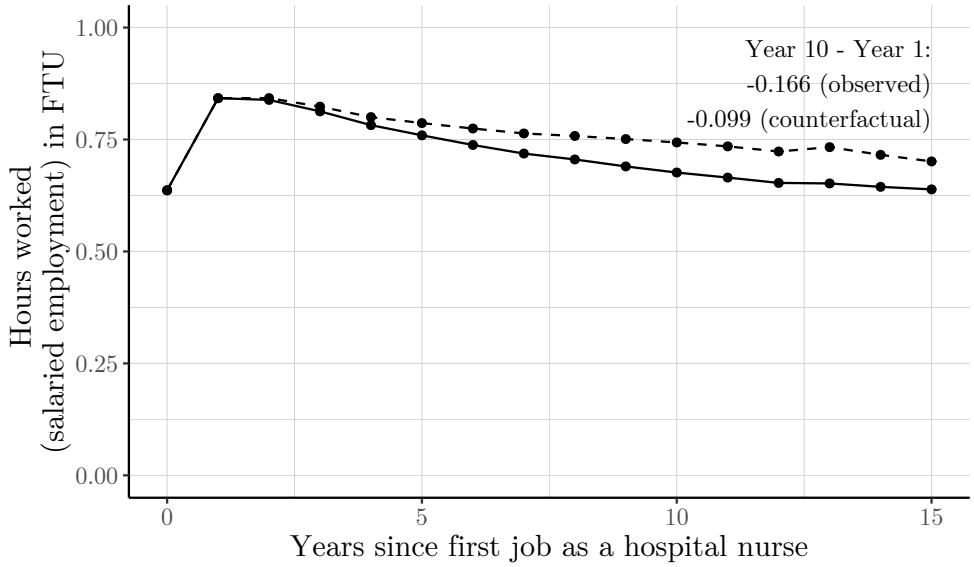

(b) Salaried employment rate

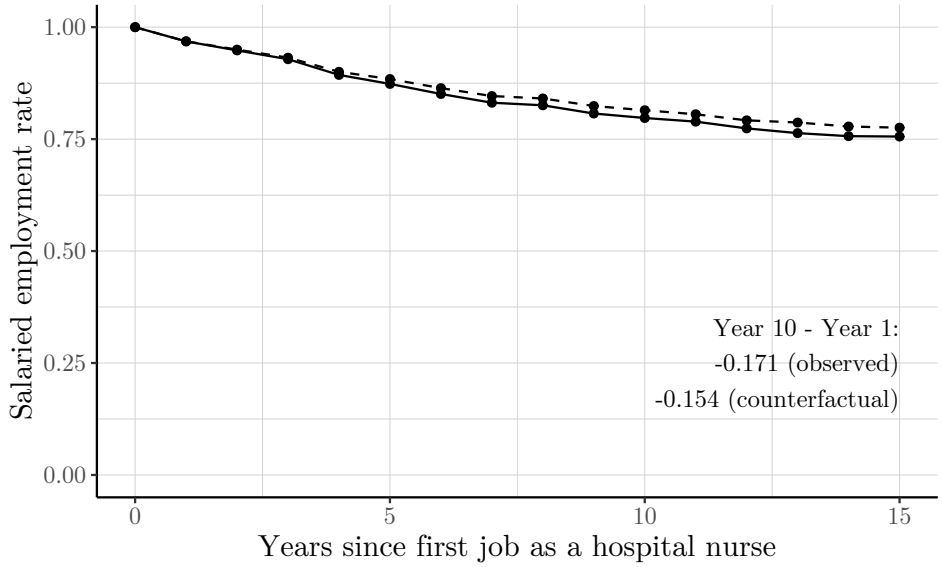

Realized and counterfactual average hours worked in the salaried sector and, in full-time units, and salaried employment rate by time relative to the first qualified healthcare worker job at a hospital. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked).

Note. Data on individuals who got their first hospital nurse job in 1988 or before are omitted from the computation.

Source. Insee, DADS-EDP panel.