Ceilings, Floors, and Children. The Gender Pay Gap over the Lifecycle*

Pierre Pora† Lionel Wilner‡

Preliminary draft, please do not quote

Abstract

Using French administrative data, we investigate the consequences of childbirth on gender inequality in the labor market between 2005 and 2015. We develop a novel approach that allows us to consider how the impact of children varies with labor market opportunities of men and women. While men are almost not affected by the arrival of a child, mothers experience large labor earnings losses that are primarily driven by labor supply decisions at both the intensive and extensive margins. Low-wages women opt out or decrease their working hours more frequently than their higher wages counterparts do. This is likely due to the opportunity cost of career interruptions, in line with incentives created by the French family insurance system. Consistent with this heterogeneous specialization towards family and the labor market, both low-wages and high achieving mothers face hourly wages penalties with respect to fathers, albeit with a different timing. As a result, the arrival of a child contributes to both glass ceilings and sticky floors patterns that are common among developed countries.

Keywords: Gender pay gap, glass ceiling, sticky floor, motherhood penalties, female labor supply.

JEL Classification: J16, J31.

*We thank Thomas Breda, Nila Ceci-Renaud, Élise Coudin, Dominique Meurs, Sébastien Roux, Milena Suarez-Castillo and Grégory Verdugo, as well as attendees at AFSE (Paris, 2018), JMS (Paris, 2018), JMA (Bordeaux, 2018) ESEM (Cologne, 2018) and Insee seminar (D2E october 2018) for useful suggestions. All errors and opinions are ours.

†Insee-CREST. Corresponding author. Address: 88 avenue Verdier, 92120 Montrouge (France). Phone: (+33)187695944. Email: pierre.pora@insee.fr

‡Insee-CREST.

1 Introduction

Despite considerable convergence in education and labor market participation over the last decades (Goldin, 2014), the gender pay gap remains a pervasive pattern in all countries. Recent research has highlighted two channels that may explain a substantial share of the persistent aggregate gender pay gap. Firstly, the earnings distribution having a fat right tail implies that massive underrepresentation of women among the top percentiles, i.e. vertical segregation can explain a very large share of the gap, that is not accounted for by traditional factors like education or occupational sorting (Fortin, Bell, and Böhm, 2017). Secondly, because genders differ in preferences over family and career, the arrival of children results in a persistent gender gap in participation, working time and hourly wages, which would account for a massive and increasing share of the aggregate gender gap (Kleven, Landais, and Søgaard, 2018). While the evidence sustaining each claim is substantial, it leads to wonder whether the consequences of childbirth are a plausible channel for women's underrepresentation among high wages workers, or more generally if children-related time constraints affect women's labor outcomes the same way given their initial level of hourly wages.

In this paper, we investigate the heterogeneity of the consequences of childbirth in terms of women's labor outcomes (total labor earnings and hourly wages) and labor supply (at both the extensive and intensive margins of employment) in the short to medium run all along the hourly wages distribution. We find that: (i) highly paid women experience much smaller labor earnings losses due to childbirth than their lower paid counterparts; (ii) they are much less likely to interrupt their careers or reduce their working hours; (iii) yet they experience somehow similar hourly wages losses; (iv) highly paid men experience slight additional hourly wages growth at the arrival of their first child so that parenthood may contribute to larger gender divergences in hourly wages growth at the top of the distribution, thus contributing to a glass ceiling pattern. We provide additional evidence that suggests these patterns stem from the opportunity cost of career interruptions and returns on experience increasing along the wages distribution (the latter contributing to the former). We also show that in the long run, children-related gender differences in hourly wages growth are larger in the lowest part of the wages distribution than they are for highly paid workers, so that they do contribute to a sticky floor effect by which poorly paid women experience slower career progressions. Lastly, we find that among individuals without child, gender differences in hourly wages growth are larger at both ends of the distribution, which suggests that glass ceilings and sticky floors are not merely the consequences of fertility.

Overall, our findings indicate that public policies aimed at increasing incentives for women to remain in employment after childbirth, relative to men could be instrumental in reducing the gender pay gap in wages progression. Indeed, since returns on experience are much larger for highly paid workers than for their lower paid counterparts, this can even be the case among the former, even though childbirth-related career interruptions and decrease in working hours are already unfrequent among them. Note that increasing incentives for women relative to men can be achieved by policies that target men, like paternity leaves. Relatedly, we believe these results draw attention to the financial incentives created by parental leave allowances. Namely, allowances like those studied by Piketty (2005); Lequien (2012); Joseph et al. (2013) provide parents that interrupt their careers to take care of their children with a fixed income that does not depend on their hourly wages. Hence, they could likely decrease the opportunity cost of career interruptions for low earning mothers with respect to that of their high achieving counterparts, thus lead some of them to get stuck between low labor force participation and low hourly wages.

Another consequence of our results is that because maternal labor supply decisions are contingent on hourly wages, the inclusion of labor market experience in wages regression, especially in the context of cross-sectional Oaxaca (1973)-Blinder (1973) decompositions of the gender pay gap, is likely to generate biased results. Specifically here it would tend to underestimate the unexplained part of the gap that can be interpreted as discrimination. While Kunze (2008) already anowledges this, we believe it is worth keeping it in mind in a context where public policies that monitor firms based on these decompositions are being implemented (see Vaccaro, 2017, for the Swiss case).

Our empirical approach combines a descriptive framework that aims at depicting heterogeneity in individual labor market trajectories along the wages distribution (Guvenen et al., 2016, 2017) with a difference-in-difference setting around childbirth similar to that of Rodrigues and Vergnat (2016). It requires large and detailed longitudinal datasets that have not always been available. We rely on the DADS-EDP panel, a French database that merges longitudinal administrative records, the filling of which is mandatory for payroll taxes, and which contains information on individual's labor earnings and paid hours, with birth and marriage records and individual census data.

The rest of the paper is organised as follows. Next section is devoted to a brief literature review. Section 3 presents our data. In section 4, we describe our empirical approach. Section 5 displays our results, and section 6 concludes.

2 Literature review

Our paper is related mostly to two strands of the literature. Firstly, the gender gap literature has now extended beyond average pay differences between men and women, in order to investigate for heterogeneity along the earnings and wages distribution, especially by relying on quantile regression (e.g. Albrecht, Björklund, and Vroman, 2003; Arulampalam, Booth, and Bryan, 2007; de la Rica, Dolado, and Llorens, 2008; Etienne and Narcy, 2010; Christofides, Polycarpou, and Vrachimis, 2013). While interpretation of these quantile regression estimates can be questionable (Gobillon, Meurs, and Roux, 2015), and estimates may vary across countries, overall results tend to converge in finding that the gap is larger among highly paid workers or at the bottom of the wages distribution than it is for median workers. These findings have been regarded as evidence for the existence of a glass ceiling – when the gap is larger at the top of the distribution – or a sticky floor – when the gap is wider at the lowest end of the distribution.

While evidence of this heterogeneity is in itself interesting, recent research has highlighted its relevance for the understanding of the average pay gap. Indeed, Fortin, Bell, and Böhm (2017) show that the huge underrepresentation of women among top earners (defined by their rank in the overall earnings distribution) explains a substantial part of the aggregate gender gap. Additionally, they argue that in a period of growing top-earnings inequality, this channel is likely to play a major part in the evolution of the gap in the future. However, it remains difficult to get precise causes for this underrepresentation, so as to implement efficient policies to act upon it. Relatedly, focusing on the very top of the earnings distribution, Guvenen, Kaplan, and Song (2014) find massive underrepresentation of women within the top percentiles. Taking advantage of the longitudinal nature of their data, they provide evidence that that the probability of moving up in the distribution is higher for men that it is for women. Reversely, women's presence at the top of the distribution is also more transitory than that of men, a stylised fact they propose to call a *paper floor*. However, because they only provide estimates for top earners, it is not clear whether women having slower earnings growth and steeper earnings losses than their male counterparts is specific to the top of the distribution, or general to all workers.

Another strand of the literature has long investigated how men and women's labor market outcomes may diverge over the lifecycle, especially because childbirths tighten time constraints and shift women's labor supply and labor market outcomes, which may thus explain a substantial share of the gender pay gap (Waldfogel, 1995, 1997, 1998). This literature includes papers that have documented how childbirths account for gender gaps in career progression among high achieving individuals, like Wood, Corcoran, and Courant (1993) or Bertrand, Goldin, and Katz (2010). It has particularly emphasized how career interruptions and decreasing labor supply of mothers lead them to poorer labor outcomes and thus contribute vastly to the increase of the gender pay gap over the lifecycle (Meurs, Pailhé, and Ponthieux, 2010). Additionally, new empirical evidence suggests that maternity may not only explain a large part of the gender gap in labor earnings, but that it may actually account for a growing share of the gap in developed countries (Kleven, Landais, and Søgaard, 2018).

Overall, authors tend to concur in finding childbirth to explain a significant share of the increase of the gender gap over the lifecycle, or more generally of the aggregate gender gap (Bertrand, Goldin, and Katz, 2010; Wilner, 2016; Adda, Dustmann, and Stevens, 2017; Juhn and McCue, 2017; Kleven, Landais, and Søgaard, 2018), though there is divergence in how much exactly, or on whether or not this contribution is increasing over time. Furthermore, evidence as to how heterogeneous can motherhood penalties be is actually quite limited, and not concurring. Anderson, Binder, and Krause (2003) find the family pay gap to be steeper among middle-skilled workers, whereas Bütikofer, Jensen, and Salvanes (2018) find it to be larger among high-skilled professions with non-linear wages structure, but Kleven, Landais, and Søgaard (2018) find very limited heterogeneity in terms of educational levels. Hence it remains difficult to assess how distributional issues pointed out by the glass ceilings and sticky floors literature are related to parenthood.

Potential endogeneity of fertility decisions has been regarded as a key empirical issue in the investigation of the labor market consequences of childbirths, leading researchers to resort to diverse approaches in order to estimate the causal effect of childbirth on labor supply or labor market outcomes: using twins birth (Rosenzweig and Wolpin, 1980), familiy background (Korenman and Neumark, 1992), siblings sex mix (Angrist and Evans, 1998) as instruments for children, or relying on propensity score matching estimates (Simonsen and Skipper, 2006). However, Kleven, Landais, and Søgaard (2018) find the causal effect of third childbirth, estimated thanks to sex-mix instruments, not to differ much from their initial estimate based on simple event-study techniques. In this paper, we rely to a large amount on this result to advocate for our rather simple difference-in-difference framework, similar to that of Rodrigues and Vergnat (2016).

Relatedly, childbirth-related labor supply decisions may also well be endogeneous with respect to subsequent labor outcomes, which advocates for cleaner identification based on exogeneous policy changes like the ones investigated by Lequien (2012); Joseph et al. (2013). In this paper, we provide evidence that childbirth-related labor supply decisions of women vary to a large extent with their hourly wages few years before the arrival of child, and are thus strongly correlated with potential future hourly wages. Using competing risk models to explain time spent out of the labor market by recent mothers, and decrease in labor supply, Rodrigues and Vergnat (2018) find similar results. This is likely to create bias in cross-sectional Oaxaca (1973)-Blinder (1973) decompositions of the gender pay gap where experience is included as a control, and may lead to underestimate the unexplained part of the gap, that is interpreted as discrimination. However, for the sake of this paper, when investigating post-childbirth career progressions, we treat these decisions as exogeneous *within groups of individuals with similar pre-birth hourly wages.*

3 Data and institutional background

3.1 The DADS-EDP panel

Our analysis is based on a large panel of French salaried employees, the longitudinal version of the *Déclarations Annuelles de Données Sociales* (DADS). By law,[1] French firms have to fill in the DADS – an annual form that is the analogue of the W-2 form in the US – for every employees affected by payroll taxes. As of year 2002, the panel contains information on individuals born on January, 2nd to 5th, April, 1st to 4th, July, 1st to 4th and October, 1st to 4th; it is therefore a representative sample of the French salaried population at rate 3,3%. Since filling in the form is mandatory, and because of the comprehensiveness of the panel with respect to individual's careers, the data is of exceptional quality and has low measurement error in comparison with survey data; it has thus this desirable feature, on top of a large sample size and no top-coding.

The database contains detailed information about gross and net wages, work days, working hours, other jobs characteristics (the beginning and the end of an employment's spell, seniority, a dummy for part-time employment), firm characteristics (industry, size, region) and inidividual characteristics (age, gender). Our variables of interest are: (i) real annual earnings defined as the sum of all salaried earnings, (ii) working time measure in working hours, and (iii) hourly wages defined as the ratio of annual earnings over working time.

Individual are identified by their NIR, a social security number with 13 digits that allows to link the DADS panel with the *Échantillon démographique permanent*, which is a longitudinal version of the censuses and census surveys and of births and marriage registers as of year 1968 for individuals born on January, 2nd to 5th, April, 1st to 4th, July, 1st to 4th and October, 1st to 4th. It thus contains information on childbirth and partial information on education. However, information on childbirth is missing before 2002 for individuals born January, Aprial or July. For this reason, we only use information on individuals born October 1st to 4th. Additionally, some part of the information related to childbirth in administrative birth registers for individuals born October 2nd to 3rd is incomplete during the 1990s, which has been documented (see Wilner, 2016), so that information has

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

to be taken from the census data rather than from birth records. In Appendix A, we explain how we proceed and provide evidence that for the sake of this paper, the quality of information for these individuals is comparable with that of individuals born October 1st or 4th for which birth records data is complete. The education variable indicates the highest degree obtained at the end of studies (see Charnoz, Coudin, and Gaini, 2011). We recode it in three categories: less than high school, high school or some college, university degree.

Our working sample is composed of male and female salaried employees working in metropolitan France between 2005 and 2015, aged 20 to 60, at the exclusion of agricultural workers and household employees.

The empirical analysis described in Section 4 requires to select individuals with a strong attachment to the labor market. Namely, we rely on "relatively stable" workers to describe their position along the wages distribution. We impose in particular that these individuals are employed in the private sector at least to years between $t-5$ and $t-2$ on top of being present in $t-1$. To deal with individuals for which labor participation is very low, we consider an individual to be employed during year t when her working hours exceed 1/8 of the annual duration of work (1607 hours as of year 2002) and her total employment duration is superior to 45 days and she earns more than 1/8 of the annual minimum wage. We also winsorize labor earnings at quantile of order 0.99999, in order to avoid issues related to potential outliers. Lastly, we drop individuals for which one observation has net labor earnings inferior (resp. superior) to 1/100 of (resp. 100 times) her gross labor earnings. In the end, our sample gathers over 1.5 million individuals-years observations, corresponding to more than 250 000 workers.

In Table 1, we give some descriptive statistics on the successive steps of the selection of "relatively stable" workers. First comes the censoring of observations with low working hours or low employment duration. Second comes the restriction to individuals that were present two year between $t-5$ and $t-2$ on top of being present in $t-1$ and t. Consistent with the rationale, both steps tend to increase average hourly wages, within gender, age groups and industry. The selection of "relatively stable" workers is harsher for women than it is for men, which is in line with them being more likely to experience career interruptions. The censoring decreases slightly the share of younger workers, which is consistent with entry in the labor workforce through shorter and non-full time employment spells. For the same reason, so does the selection of "relatively stable" workers. The censoring decreases the share of workers in the service industry, which is in line with them being more likely to have short employment spells and part-time employment. The selection of "relatively stable" workers also decreases the share of service industry workers among men, and the share of trade industry workers among women. This may result from service industry male workers (resp. trade industry female workers) having more unstable employment histories than their counterparts working in other industries.

Both within our base sample, after the censoring and among "relatively stable" workers, the gender gap in hourly wages is larger among older workers than it is among their younger counterparts.

Table 1 – Descriptive statistics on the selection process NOT UP TO DATE

	Base sample				Censoring				Final sample			
N_{it}	9 940 285				9 165 771				6 823 722			
N_i	1 243 794				1 158 063				944 387			
	Women		Men		Women		Men		Women		Men	
N_{it}	4 369 060		5 571 225		3 910 929		5 254 842		2 843 890		3 979 832	
N_i	581 663		682 131		526 979		631 084		414 732		529 655	
	Frequency (in %)	Average hourly wages (2015 €)	Frequency (in %)	Average hourly wages (2015 €)	Frequency (in %)	Average hourly wages (2015 €)	Frequency (in %)	Average hourly wages (2015 €)	Frequency (in %)	Average hourly wages (2015 €)	Frequency (in %)	Average hourly wages (2015 €)
Age												
23-24	5,9	9,46	6,9	10,03	6,4	9,51	5,5	10,10	3,7	9,45	3,5	10,44
25-29	15,4	10,83	16,7	11,92	16,4	10,97	15,1	12,03	14,8	11,18	13,5	12,37
30-34	16,0	12,09	15,8	13,92	15,8	12,28	16,0	14,05	15,4	12,85	15,3	14,55
35-39	15,5	12,68	14,8	15,34	14,9	12,88	15,6	15,48	14,8	13,46	15,6	15,81
40-44	14,6	12,75	14,1	16,98	14,3	12,92	14,7	17,15	14,5	13,41	15,2	17,22
45-49	12,9	13,24	12,7	17,37	13,0	13,44	13,1	17,53	13,6	14,07	13,7	18,00
50-54	11,3	13,00	10,9	18,12	11,2	13,18	11,5	18,3	12,0	13,48	12,1	18,73
55-59	8,4	13,33	8,0	18,78	8,1	13,52	8,4	19,00	11,0	13,99	11,0	19,70
Industry												
Manufacturing	12,9	12,93	24,6	15,9	13,9	13,05	25,7	15,98	14,4	13,9	26,7	16,41
Construction	1,9	12,77	11,9	13,38	2	12,85	12,3	13,41	2,1	13,84	12,5	14,26
Trade	19,7	10,76	15,5	14,7	20,2	10,88	15,8	14,89	20,1	11,6	15,9	15,56
Services	65,6	12,07	48	14,96	63,9	12,36	46,2	15,37	63,3	13,19	44,9	16,57

3.2 Institutional background

Family-friendly policies in France have a long-lasting history (see Rosental, 2010) that dates back at least from pro-natalist concerns in the interwar period (Huss, 1990). These policies rely on (i) tax cuts, especially the *quotient familial* introduced in 1945 by which the income tax rate depends on the number of children in a household; (ii) various child benefits; (iii) some other welfare benefits, such as bonuses for retirement pensions that depends on realized fertility, or housing allowances. Note that in France income in taxed jointly within households, which for low incomes may create strong incentives towards within-household specialization (Carbonnier, 2014).

Maternity leaves were created in 1909, first being unpaid, and are fully covered by social insurance up to a threshold for all salaried workers since 1970. Since 1980, the arrival of the first two children grants women with a 16 weeks maternity leave, 6 weeks before childbirth and 10 weeks after. From the arrival of the third child, the total duration is 26 weeks (8+18), and maternity leave duration may go up to 46 weeks in the case of multiple births. Maternity leaves also come with a minimum duration of 8 weeks, that is 2 weeks before childbirth and 6 weeks after.

Paternity leaves were enforced in 2002. It grants fathers with a 11 days-long leave that is fully covered by social insurance up to a threshold. The duration can go up to 18 days in the case of multiple births, and it includes weekends and public holidays. This paternity leave is granted on top of birth leaves that amount to 3 days.

On top of these maternity and paternity leaves come various parental allowances. Namely, the PAJE (*Prestation d'accueil du jeune enfant*), created in 2004 comes as a package. It comprises a one-shot means-tested bonus at childbirth (*prime de naissance*), monthly means-tested benefits (*allocations familiales*), a childcare subsidy (*complément libre choix du mode de garde* or CMG), and some child benefit that are granted when parents interrupt their careers or work part-time.

These benefits date back to 1985 with the creation of the APE (*Allocation Parentale d'Éducation*) that was initially restricted to mothers of 3 or more. In 1994, the APE was extended to mothers of 2 (Choné, Le Blanc, and Robert-Bobée, 2004; Piketty, 2005; Lequien, 2012), and was replaced in 2004 by the CLCA, from the 1st child onwards (Joseph et al., 2013), which provides a fixed amount that is not mean-tested, for a maximal duration of 6 months. Lastly, in 2015 the CLCA was replaced by PreParE (*Prestation partagée d'éducation de l'enfant*) to which fathers become eligible.

Other policies favor participation to the labor force by decreasing the cost of childcare, like the CMG, which is not means-tested (Givord and Marbot, 2015), or tax credits up to 50%.

4 Empirical analysis

Our main outcome of interest is total annual labor earnings of individual i during year t, that we denote as $\tilde{y}_{it}$. We decompose it into four components: d_{it} a dummy for participation, $\tilde{x}_{it}$ the employment duration in days, comprised between 0 and 360, $\tilde{h}_{it}$ the average number of working hours per day during year t, and lastly $\tilde{w}_{it}$ the average hourly wages of individual i during year t. Hence:

$$\tilde{y}_{it} = d_{it}\tilde{x}_{it}\tilde{h}_{it}\tilde{w}_{it} \tag{1}$$

4.1 Normalization

The first step of our empirical framework derived from that of Guvenen et al. (2016, 2017) is normalizing earnings and its components with respect to age, cohort and period. Let $\tilde{z}$ denote either labor earnings or one of its component, with the exception of the participation dummy. We start by regressing (log of) $\tilde{z}_{it}$ on a full set of cohort (year of birth), age and period dummies. We do so on the full sample of individuals aged 20 to 60 that are employed in the private sector at time t. Recall that participants have total working hours superior to 1/8 of the annual legal duration of work, total employment duration superior to 45 days and total labor earnings superior to 1/8 of the annual minimum wage. We estimate this regression as a pooled cross-section:

$$\ln(\tilde{z}_{it}) = \sum_c \lambda_c^z \mathbb{1}_{cohort_i=c} + \sum_a \mu_a^z \mathbb{1}_{age_{it}=a} + \sum_T \nu_T^z \mathbb{1}_{t=T} + \epsilon_{it}^z \tag{2}$$

The identification of age-period-cohort (APC) models can be achieved at the cost of some normalizations. The major threat to the simultaneous identification of λ, μ and ν stems from colinearity between age, cohort and period: age is equal to current period minus year-of-birth. Several solutions have been investigated in the sociological literature, e.g. Mason et al. (1973) who propose to assume that any two ages, periods or cohort have the same effect, on top of removing one dummy in each dimension. Deaton and Paxson (1994) and Deaton (1997) suggest a transformation[2] of period effects in order to meet two requirements: (i) these time effects sum to zero, and (ii) they are orthogonal to a time trend, so that

[2] An insightful presentation of this method is provided by Afsa and Buffeteau (2006).

age and cohort effects capture growth while year dummies account for cyclical fluctuations or business cycle effects that average to zero over the long run. To sum up, the parameters of the model (λ, μ, ν) are identified provided that $\lambda_{\underline{c}} = 0$ and $\sum_{t=1}^{T} \nu_t (t-1) = 0$. The corresponding transformation of time dummies $d_T = \mathbb{1}_{t=T}$ writes as follows:

$$d_T^* = d_T - [(T-1)d_2 - (T-2)d_1] \tag{3}$$

with $d_1^* = d_2^* = 0$. In practice, it is convenient to include all age dummies, all cohort dummies but the first and all transformed dummies d_T^* but d_1^* and d_2^* in the regression.

Accounting decomposition 1 implies that $\lambda_c^y = \lambda_c^x + \lambda_c^h + \lambda_c^w$ and similar identites for age and period coefficients. We use our estimates to define the normalized component z_{it} as:

$$z_{it} = \frac{\tilde{z}_{it}}{\exp(\hat{\lambda}_{cohort_i} + \hat{\mu}_{age_{it}} + \hat{\nu}_t)} \tag{4}$$

An accounting decomposition similar as 1 stands for normalized earnings:

$$y_{it} = d_{it} x_{it} h_{it} w_{it} \tag{5}$$

4.2 Ranks in the hourly wages distribution

Our empirical strategy interacts a difference-in-difference setting with a descriptive framework that aims at depicting heterogeneity in the consequences of childbirth along the wages distribution. In order to do so, we rely on comparisons *within* groups of workers that have similar hourly wages – this is the difference-in-difference aspect of our approach – and comparisons of these estimates *across* groups of workers characterized by their hourly wages. Hence the definition of these groups is key to our analysis. It is based on a measure of recent hourly wages:

$$W_{it} = \frac{\sum_{T=t-5}^{t-1} d_{iT} \tilde{w}_{iT}}{\sum_{T=t-5}^{t-1} d_{iT} \exp(\hat{\lambda}_{cohort_i} + \hat{\mu}_{age_{iT}} + \hat{\nu}_T)} \tag{6}$$

We only compute this measure for individuals that participate in $t-1$ and at least twice between $t-5$ and $t-2$, that is provided that $d_{i,t-1} \sum_{T=t-5}^{t-1} d_{iT} \geq 3$.

Lastly, within each age × year cell, we rank workers according to their recent wages W_{it}. We use this ranking to create 20 cells: P0-P5, P5-P10, ..., P90-P95 and P95-P100. Hence within each age × year × recent wages cell, we can consider workers that are if not identical, at least *ex ante* quite similar with respect to their hourly wages levels before year t. It is worth noting that ranks are not conditional on gender so that within these cells men and women have approximately the same recent wages.

It is worth noting that this depiction of heterogeneity along the wages distribution yields conceptually different estimates than a conditional quantile regression based approach Koenker and Bassett (1978) would. A conditional quantile regression based estimation gives information as to how heterogeneous the effect of a given covariate on outcome is, but does not relate it to any specific dimension that we would be interested in. In the case where individuals fixed effects are introducted, relying on conditional quantile regression may give quantile treatment effects, i.e. the distribution of changes in individuals' outcomes in response to a change in the covariate, but there is no reason why the upper quantiles of these changes would coincide for instance with individuals that belong the the highest end of the outcomes distribution. Firpo, Fortin, and Lemieux (2009) develop an unconditional quantile regression framework that allows easier interpretation of the coefficients in terms of the outcomes distribution, but it does not seem well-suited to the case of longitudinal data. Recent attemps have nevertheless been made to expand unconditional quantile regression to panel data (Powell, 2016), but we stick to a much simpler approach where we focus on average treatment effects, solely interacting our categorical depiction of the recent wages distribution with our covariates of interest.

4.3 Difference-in-difference setting

Our estimates of the consequences of childbirth are based on a difference-in-difference framework. We define our treatment as experiencing childbirth during some year t. Because our data based on birth-records and censuses gives information on the rank of childbirth, we use different control groups depending on whether we consider first childbirth, second childbirth etc. Specifically, similar to Rodrigues and Vergnat (2016), our control group for first childbirth is composed of individuals (of the same gender) who never had children ; our control group for second childbirth is composed of individuals with one child who never had a second child etc. Note that the right censoring of the data in 2015 creates measurement errors, so that individuals of our nth control group may have nth child born after 2015. In practice, we limit ourselves to the first three childbirths, that represent 96% of childbirths in our data.

This difference-in-difference approach is embedded in our ranking based on recent wages. As a consequence, our control groups for childbirth are based on the above definition, and are also restricted to individuals that had the same rank in the recent hourly wages distribution than our treated individuals. What is more, the effect of childbirth is also allowed to vary along the recent wages distribution. Lastly, because our framework imposes that individuals of our sample participate in the labor market during year $t-1$, we take this year as a reference.

4.3.1 Descriptive approach

We first rely on a descriptive framework to provide estimates of childbirth on labor market outcomes and labor supply. Our estimate of the consequences of nth childbirth on earnings k years after childbirth for individuals of gender g at rank r in the recent wages distribution writes:

$$\begin{aligned}\beta_{g,r}^{y,n,k} = & \underbrace{\ln(\mathbb{E}[y_{i,t+k}|b_{it}^n = 1, r_{it} = r, g_i = g]) - \ln(\mathbb{E}[y_{i,t-1}|b_{it}^n = 1, r_{it} = r, g_i = g])}_{\text{Treated}} \\ & - \underbrace{\ln(\mathbb{E}[y_{i,t+k}|c_{it}^n = 1, r_{it} = r, g_i = g]) - \ln(\mathbb{E}[y_{i,t-1}|c_{it}^n = 1, r_{it} = r, g_i = g])}_{\text{Control}} \end{aligned} \quad (7)$$

where b_{it}^n is a dummy for experiencing nth childbirth during year t and c_{it}^n is a dummy for belonging to the nth control group at time t, i.e. having $n-1$ children at time t but never experiencing nth childbirth according to the data.

Accounting decomposition 5 implies that the logarithm of average normalized earnings growth can be decomposed in a sum of its four components plus a selection term that arises from the fact that individuals that participate in $t+k$ may not have the exact same past earnings $y_{i,t-1}$ as those who do not participate:

$$\ln\left(\frac{\mathbb{E}\left[\frac{y_{i,t+k}}{y_{i,t-1}}y_{i,t-1}\right]}{\mathbb{E}[y_{i,t-1}]}\right) = \underbrace{\ln\left(\frac{\mathbb{E}[y_{i,t-1}|d_{i,t+k}=1]}{\mathbb{E}[y_{i,t-1}]}\right)}_{\text{Selection}} + \underbrace{\ln\left(\mathbb{P}(d_{i,t+k}=1)\right)}_{\text{Participation}} + \underbrace{\ln\left(\frac{\mathbb{E}\left[\frac{x_{i,t+k}}{x_{i,t-1}}x_{i,t-1}h_{i,t-1}w_{i,t-1}|d_{i,t+k}=1\right]}{\mathbb{E}\left[x_{i,t-1}h_{i,t-1}w_{i,t-1}|d_{i,t+k}=1\right]}\right)}_{\text{Employment Duration Changes}} + \underbrace{\ln\left(\frac{\mathbb{E}\left[\frac{h_{i,t+k}}{h_{i,t-1}}x_{i,t+k}h_{i,t-1}w_{i,t-1}|d_{i,t+k}=1\right]}{\mathbb{E}\left[x_{i,t+k}h_{i,t-1}w_{i,t-1}|d_{i,t+k}=1\right]}\right)}_{\text{Hours Per Day Changes}} + \underbrace{\ln\left(\frac{\mathbb{E}\left[\frac{w_{i,t+k}}{w_{i,t-1}}x_{i,t+k}h_{i,t+k}w_{i,t-1}|d_{i,t+k}=1\right]}{\mathbb{E}\left[x_{i,t+k}h_{i,t+k}w_{i,t-1}|d_{i,t+k}=1\right]}\right)}_{\text{Hourly Wages Growth}} \tag{8}$$

This accounting decomposition of labor earnings growth allows us to separate each component of the consequences of childbirth on earnings: $\beta^{y,n,k}_{g,r} = \beta^{s,n,k}_{g,r} + \beta^{d,n,k}_{g,r} + \beta^{x,n,k}_{g,r} + \beta^{h,n,k}_{g,r} + \beta^{w,n,k}_{g,r}$ where $\beta^{s,n,k}_{g,r}$ stand for the selection term and the four others correspond to each component of labor earnings.

An interesting feature of our method and our data is that by taking $k < -1$ we can the exact same way test the parallel trend assumption upon which difference-in-difference estimators are based.

4.3.2 Regression framework

We also implement related regressions in order to estimate how much childbirth affects the gender gap in career progression. These regressions also enable us to consider additional covariates, in order to get a better understanding of the channels that drive our descriptive results. Here, our estimate of the consequences of nth childbirth on earnings k years after childbirth for individuals of gender g at rank r in the recent wages distribution writes:

$$\begin{aligned}\beta_{g,r}^{y,n,k} = &\underbrace{\mathbb{E}[\ln(y_{i,t+k}) - \ln(y_{i,t-1}) | b_{it}^n = 1, r_{it} = r, g_i = g]}_{\text{Treated}} \\ &- \underbrace{\mathbb{E}[\ln(y_{i,t+k}) - \ln(y_{i,t-1}) | c_{it}^n = 1, r_{it} = r, g_i = g]}_{\text{Control}} \end{aligned} \tag{9}$$

where b_{it}^n is a dummy for experiencing nth childbirth during year t and c_{it}^n is a dummy for belonging to the nth control group at time t.

Given a component z (with the exception of the participation dummy) of the accounting decomposition of earnings 5, we consider $\delta^k z_{i,t} = \ln(z_{i,t+k}) - \ln(z_{i,t-1})$ the growth in component z for individual t between $t-1$ and $t+k$, which is defined for all individuals that participate in the private sector during year $t+k$. We take this growth as the outcome of our regression. When considering participation, we use the participation dummy $d_{i,t+k}$ as the outcome. Because all observations have $d_{i,t-1}$ this will indeed capture changes in labor supply at the extensive margin. We estimate a $(k+1)$-difference regression by OLS:

$$\begin{aligned}\delta^k z_{i,t} = \alpha_{g_i,r_{it}}^{z,k} + \sum_n \gamma_{g_i,r_{it}}^{z,n,k}(b_{it}^n + c_{it}^n) + \sum_n \beta_{g_i,r_{it}}^{z,n,k} b_{it}^n \\ + \zeta_{g_i,r_{it}}^{z,k} X_{i,t} + u_{i,t}\end{aligned} \tag{10}$$

where X_{it} is a vector of either invariant or time-varying covariates, and u_{it} and idiosyncratic error of mean 0.

Treatment variables b_{it}^n intervene twice in 10: once summed with control dummies c_{it}^n and multiplied par parameters $\gamma_{g,r}^{z,n,k}$, and once alone and multiplied by parameters $\beta_{g,r}^{z,n,k}$. Hence parameters $\beta_{g,r}^{z,n,k}$ will capture how treated individuals change with respect to those of the treated group between $t-1$ and $t+k$, which makes $\hat{\beta}_{g,r}^{z,n,k}$ difference-in-difference estimators. In other words, the $\beta_{g,r}^{z,n,k}$ tell us how, k years after childbirth, parents' outcomes change with respect to those of non-parents of the same gender that had similar hourly wages. The main interest of our approach is that being indexed by r that denotes the rank in the recent wages distribution (conditional on age and year), the impact of childbirth is allowed to vary in a non-parametric way all along the recent wages distribution.

An alternate specification of the same regression is:

$$
\begin{aligned}
\delta^k z_{i,t} = & (\alpha^{z,k}_{r_{it}} + \alpha^{z,k}_{gap,r_{it}} g_i) \\
& + \sum_n (\gamma^{z,n,k}_{r_{it}} + \gamma^{z,n,k}_{gap,r_{it}} g_i)(b^n_{it} + c^n_{it}) \\
& + \sum_n (\beta^{z,n,k}_{r_{it}} + \beta^{z,n,k}_{gap,r_{it}} g_i) b^n_{it} \\
& + (\zeta^{z,k}_{r_{it}} X_{i,t} + \zeta^{z,k}_{gap,r_{it}} X_{i,t} g_i) X_{i,t} \\
& + u_{i,t}
\end{aligned}
\tag{11}
$$

Here our coefficients of interest are again $\beta^{z,n,k}_r$ and $\beta^{z,n,k}_{gap,r}$. The first one correspond to the impact of childbirth on father's labor outcomes. The second ones give us information as to how mother's outcomes shift with respect to those of fathers that previously earned similar hourly wages, having already controled for the divergence that occur between gender among individuals that never had children according to our data. Hence it gives us a sense of how childbirth is likely to contribute to the gender gap in career progression. Hence the $\beta^{z,n,k}_{gap,r}$ can be thought of like triple-difference estimators.

A same individual will intervene several times in these regressions. What is more, when $k > 1$ their will be some overlapping between observations related to the same individual. This will create serial autocorrelation in the error terms. Proper inference has to take this issue into account. In order to do so, we cluster standard error at the individual level (Bertrand, Duflo, and Mullainathan, 2004).

Lastly, we want to understand what drives the heterogeneity in the $\beta^{z,n,k}_{g,r}$. Namely, individuals with different initial levels in the wages rate may react differently to childbirth for two reasons. Firstly, people will make their childbirth-related labor supply decisions based on their potential wages rate, which in the short to medium run is likely to remain close to the hourly wages they earned shortly before childbirth. Here the wages rate acts directly like an incentive to remain or not in the labor workforce, and to reduce or increase working hours. Secondly, individuals are likely to differ in their unobserved preference over family and career; hence, as they anticipate more time spent out of the labor market, those with higher taste for family over career will invest less in the acquisition of labor market valued skills, and therefore earn lower hourly wages prior to childbirth (Becker, 1981). Additionaly, they may choose to work in firms that are less demanding in terms of family-career conciliation, and pay lower wages (Card, Cardoso, and Kline, 2016; Coudin, Maillard, and To, 2018). This second channel leads to a reverse causality bias that forbids to interpret the $\beta_{g,r}^{z,n,k}$ as reflecting the direct effect of hourly wages on childbirth consequences.

While we lack exogeneous variations that would enable us to provide clean identification of the direct effect of hourly wages on childbirth-related labor supply decisions and labor outcomes changes, we propose to assess in a simple way how much the second channel is likely to explain our results. In order to do so, we interact our difference-in-difference regression not only with rank in the recent wages distribution like in 10, but also with other variables that we consider to capture at least some part of past labor market valued human capital investment and preferences over family and career: education, measured by the possession of a university degree, working full-time in $t-1$, and the share of female working part-time in the firm where individual i works in $t-1$. To put it differently, we estimate:

$$
\begin{aligned}
\delta^k z_{i,t} = &(\bar{\alpha}_{g_i,r_{it}}^{z,k} + \bar{\alpha}_{g_i,Z}^{z,k} Z_{i,t}) \\
&+ \sum_n (\bar{\gamma}_{g_i,r_{it}}^{z,n,k} + \bar{\gamma}_{g_i,Z}^{z,n,k} Z_{i,t})(b_{it}^n + c_{it}^n) \\
&+ \sum_n (\bar{\beta}_{g_i,r_{it}}^{z,n,k} + \bar{\beta}_{g_i,Z}^{z,n,k} Z_{i,t}) b_{it}^n \\
&+ (\bar{\zeta}_{g_i,r_{it}}^{z,k} + \bar{\zeta}_{g_i,Z}^{z,k} Z_{i,t}) X_{i,t} + u_{i,t}
\end{aligned} \tag{12}
$$

Here heterogeneity in the $\bar{\beta}_{g,r}^{z,k}$ will stem from variations along the wages distribution of childbirth-related changes in z, *within* groups of individuals with similar Z, that is within groups of individuals with similar education, labor force attachment and firm composition as measured in $t-1$. This is certainly not sufficient to capture all the variation that arises from different human capital investment and sorting due to heterogeneous preferences over family and career. However, we believe finding substantial heterogeneity in the $\bar{\beta}_{g,r}^{z,k}$ can be considered as suggestive evidence that contemporary hourly wages is an incentive that drives much of childbirth-related labor decisions.

5 Results

5.1 Heterogeneous consequences of childbirth

We first assess the consequences of childbirth on men's and women's labor outcomes, relying on a descriptive version of a difference-in-difference. We plot our estimates of the impact of the first three childbirths on individuals' total labor earnings for men (Figure 1) and for men (Figure 2), and their decomposition into participation, days of work, working hours per day and hourly wages, plus a selection term. Treatment is experiencing nth childbirth during year t, and our control group is composed of individuals that have $n-1$ children at time t but never have an nth child in the data, that is until 2015. Our empirical framework allows us to consider how the impact of the treatment on labor earnings varies all along the hourly wages distribution defined by rank in the distribution of recent wages. We plot our estimate for $t+k \in \{t-3, ..., t+5\}$ with the exception of $t-1$; because $t-1$ is taken as a reference our estimates are all equal to 0 for this date.

Mothers experience large earnings losses after childbirth relative to women that earned similar hourly wages few years before. All components participate to these losses: after the arrival of a child, mothers are more likely to leave employment, work fewer days, work fewer hours per day and earn lower hourly wages than women of our control groups. In the short to medium run nevertheless, labor supply decisions seem to be driving these large earnings losses, while the wages rate accounts only for a limited fraction of them. What is more, the consequences of childbirth on women's labor outcomes appear to increase with rank of the child, i.e. they are harsher for the arrival of second or third child than they are for the arrival of first child.

The main point of Figure 1 is to show that children-related earnings losses among women display vast heterogeneity: women that before childbirth ranked at the bottom of the hourly wages distribution experience far larger earnings losses than their counterparts that earned higher hourly wages. Our estimates suggest that at the very bottom of the distribution, women's losses amount to 70 log-points the year they first give birth, 45 log-points one year after childbirth, and 50 log-points 5 years after the arrival of a child. By contrast, women ranked in the top 5% of the houly wages distributon before their first childbirth would experience losses that amount to 20 log-points the year they give birth, 5 log-points one year later and 0 five years after the arrival of a child.

This massive heterogeneity is primarily driven by labor supply decisions at the extensive margin: childbirth would reduce by 20 log-points (resp. 70 and 85 log-points) the probability that women are employed one year after the arrival of their first (resp. second and third) child at the lowest end of the hourly wages distribution, but would actually not decrease this probability for women that belonged to the top 5% our the hourly wages distribution. Conversely, while hourly wages losses display a U-shape pattern along the wages distribution the year of the arrival of a child, one to five years later those motherhood wages penalties look much more homogeneous and amount to approximately 5 log-points for first child (and even less for second and third children). Moreover, this U-shape pattern is likely to arise from the institutional setting of maternity leave compensation that involves various thresholds and depends on its duration, and should therefore not be taken at face value.

A nice feature of our descriptive approach is that it allows a graphical test for the parallel trend assumption upon which the difference-in-difference setting rests. Here this assumption states that the difference between our treated and control groups before $t-1$ should be equal to 0. Here this assumption is rejected by the data: there are small differences between our treated and control groups' earnings in $t-3$ and $t-2$ relative to earnings at time $t-1$. The difference is slightly positive (resp. negative) when considering the arrival of first (resp. second) child, which means that mothers had slightly slower (resp. faster) earnings growth than non-mothers (resp. mothers of one) prior to first (resp. second) childbirth. However, these differences are small (less than 10 log-points) with respect to those we measure after childbirth (up to 130 log-points). Furthemore, we find only slight heterogeneity in these differences. Based on this, we argue that even though this is a threat to the interpretation of our results as exact point estimates of the causal effect of childbirth, they still provide correct identification of the heterogeneity of the consequences of childbirth on women's labor outcomes, which is our main focus.

When it comes to men, our estimates suggest childbirth may tend to increase slightly labor earnings after the arrival of first child, especially through increased participation over the years that follow the arrival of first child, and increased hourly wages. Increased participation might be slightly stronger for fathers that are ranked at the top of the recent wages distribution.

On top of these results that provide average treatment effects, in Appendix B we provide with additional figures where we display results based on median changes as opposed to average changes, in order to check that our results do not stem from unfrequent yet very large changes. Overall we find that our results are somehow robust in the sense that they do not arise from rare yet considerable events.

5.2 What drives the heterogeneity in mother's labor supply decisions ?

With this evidence in mind, we dig further in the data to understand what drives those differences in labor supply decisions at the extensive margin between mothers who earn low wages prior to childbirth and are likely to interrupt their career after the arrival of a child, and their counterparts that earn high hourly wages and a unlikely to leave employment. Two channels are likely to contribute to these differences. Firstly, if individuals make their labor supply decisions based on the relative value of one hour spent in employment and one hour spent outside of employment, hourly wages will enter at the first order in the expression of the value of the former. Therefore, the value of one hour spent in employment is likely to be higher for highly paid individuals relative to one hour spent in home-production than it is for poorly paid individuals. This opportunity cost channel would therefore make mothers that previously earned high wages much less likely to interrupt their career than those who ranked at the bottom of the wages distribution.

Secondly, mothers that belong to different part of the wages distribution may differ with respect to their unobseved preferences towards family and career. Women with higher taste for family over career are more likely to take time outside the labor workforce after the arrival of a child than those with higher taste for career. If they are forward-looking, prior to childbirth, the former would therefore be less commited to their jobs, and invest less in the acquisition of labor market valued skills than the latter, because they anticipate more time spent outside

Figure 1 – Consequences of childbirth on women's labor outcomes: by rank in the recent wages distribution and rank of child

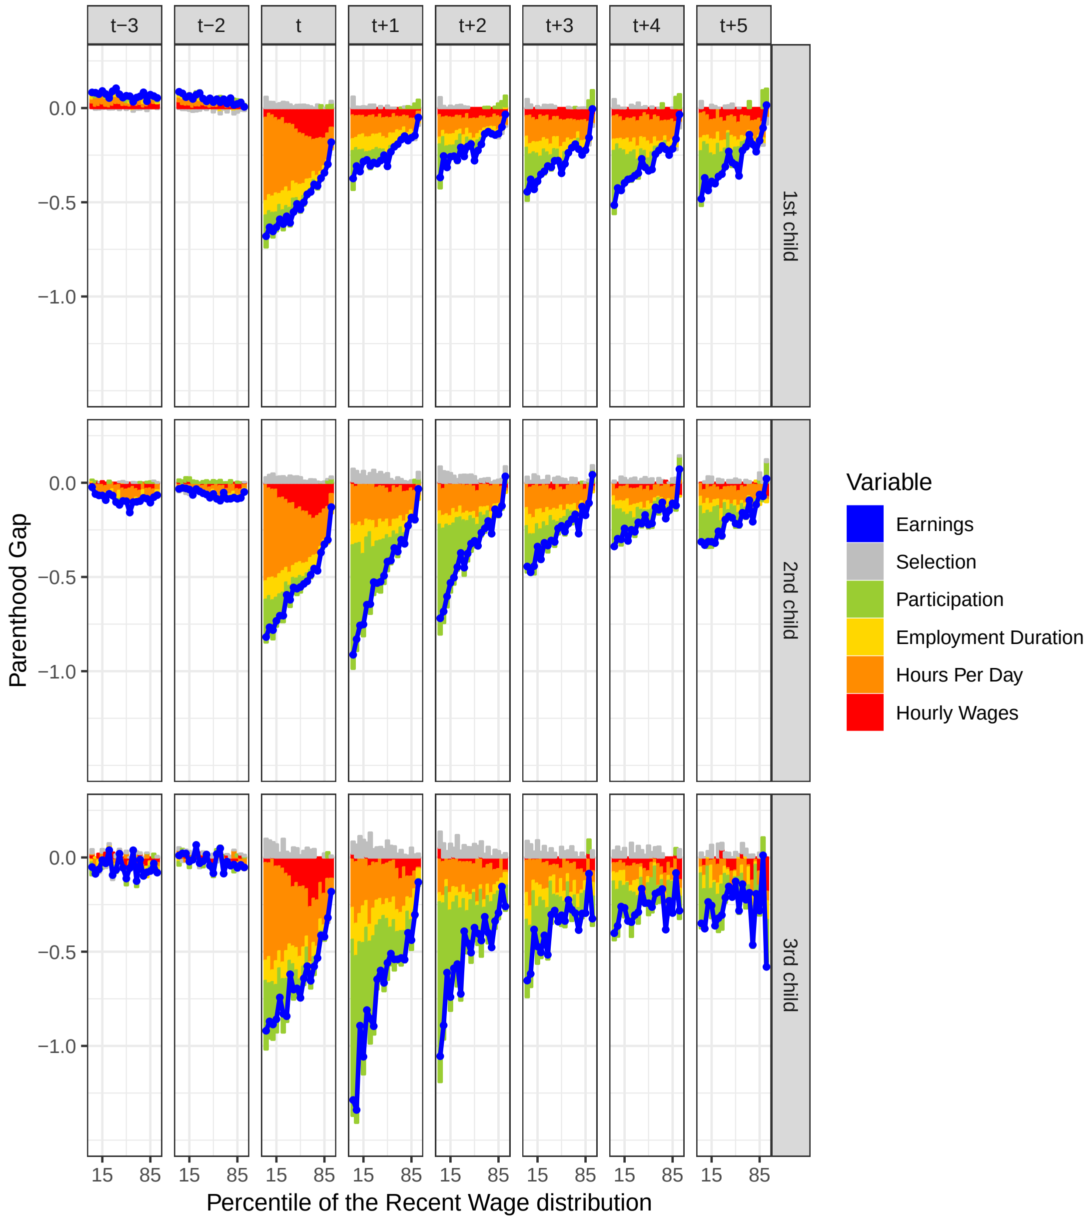

Figure 2 – Consequences of childbirth on men's labor outcomes: by rank in the recent wages distribution and rank of child

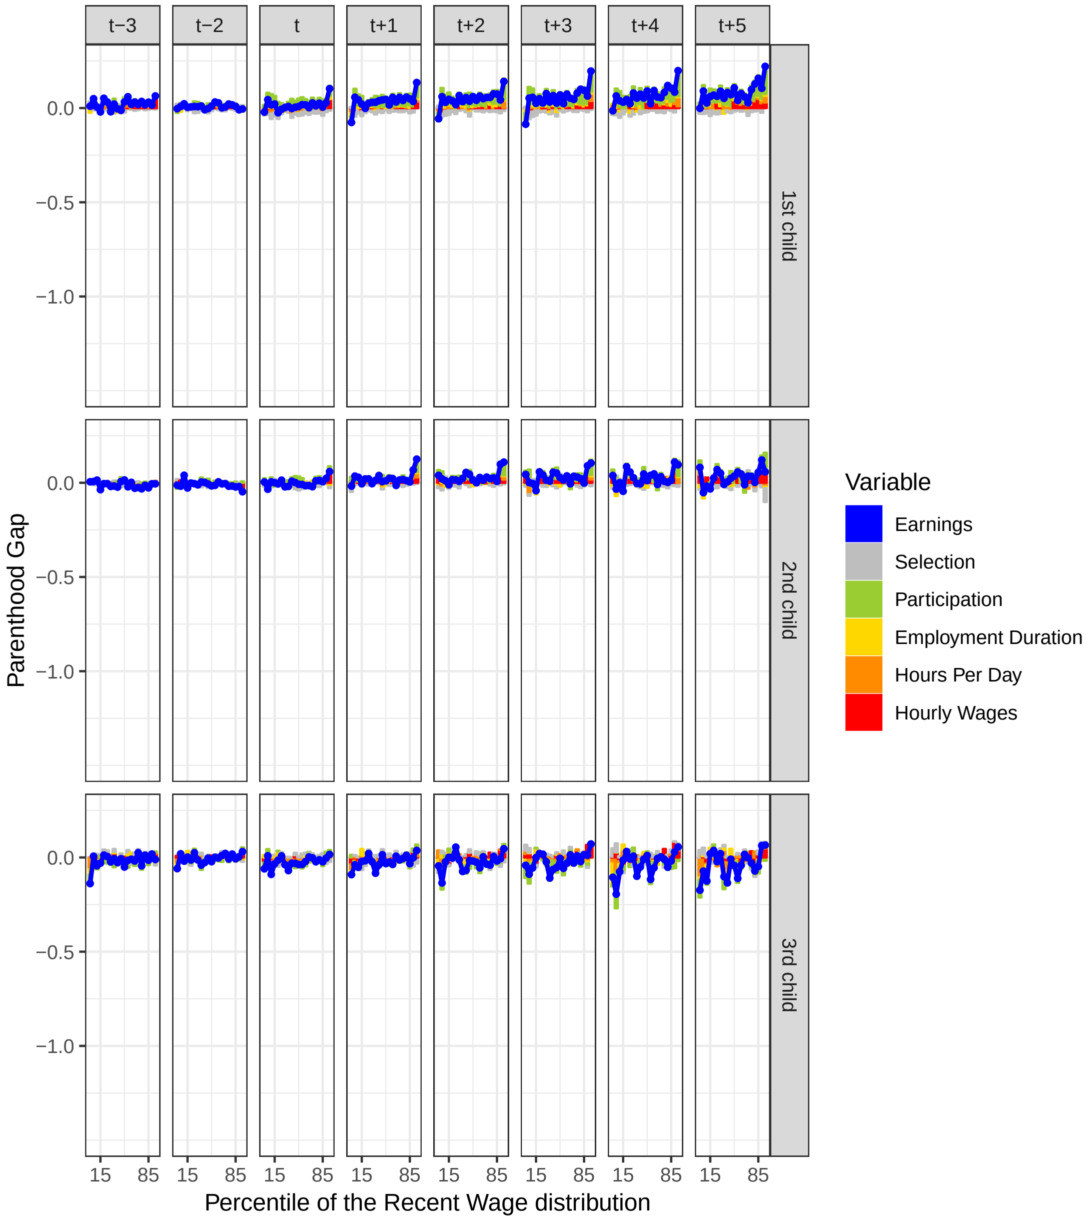

employment in the future. This difference can be rationalized in a model with two types of human capital, one that is is valued on the labor market, and one that is valued in home production (Becker, 1981). Women with higher preferences for family over career could also choose to work in firms that are less demanding in terms of work-family conciliation, and pay lower wages, which would be consistent with the view that differential sorting between genders explains may explain a significant share of the gender pay gap (Card, Cardoso, and Kline, 2016; Coudin, Maillard, and To, 2018) This undercommitment and underinvestment in labor market valued skills, and sorting between firms based on preferences over family and career will translate in lower wages prior to childbirth, and therefore in heterogeneity in labor supply response to childbirth along the hourly wages distribution.

Disentangling those two channels requires exogeneous shocks on hourly wages, that would affect the labor supply decisions related to childbirth while being independent of human capital investment decisions and sorting between firms made before the arrival of a child. We lack this source of identification in our setting. However, we propose to approximate the contribution of these channels by relying on two proxies of human capital investment and job commitment prior to childbirth: having a university degree, and working full-time the year before childbirth. We also rely on a proxy for between-firm differences in work-family conciliation: the share of women working part-time in among employees working in the same firm as individual i at time $t-1$. These proxies cannot capture all variation related to human capital investiment, job commitment and differential sorting, but we believe they should catch a substantial part of it. Our approach is then to compare the heterogeneity in children-related labor supply decisions between two settings: one that is the regression counterpart of Figures 1 and 2, where labor supply decisions can only vary with rank in the hourly wages distribution, and one where decisions can also vary depending on education, full-time vs. part-time status one year before childbirth, and share of women working part-time in firm at which an individual is employed in $t-1$. In the second setting, heterogeneity between mothers with different hourly wages prior to childbirth will stem from differences between women with similar education, full-time vs. part-time status in $t-1$ and firm-composition, so that we expect those differences to be smaller than in the first setting. The exact extent by which allowing for education and previous full-time status reduces those differences will indicate how likely it is that this heterogeneity is related to the opportunity cost channel as opposed to the unobserved preferences channel.

We operationalize this approach by considering the probability of remaining in employment one year after childbirth. We estimate a linear probability model that interact our difference-in-difference setting with our descriptive framework. Figure 3 displays the coefficients that depict heterogeneity along the recent wages distribution, first in the case where labor supply decisions can only vary depending on recent hourly wages (first and second panels), and then in the case where they can also differ depending on education, previous full-time vs. part-time status and firm composition (third and fourth panels).

First, consistent with Figure 1, when solely interacting our difference-in-difference setting with the recent wages distribution, we find that women that previously earned low hourly wages are far more likely to interrupt their careers after the arrival of a child than those that earned high hourly wages. Our results actually suggest a slightly positive effect of childbirth on labor supply a the top of the hourly wages distribution. The difference between those at the lower end of the distribution and those at the top amounts to 18 probability points for first chilbirth, 33 points for second childbirth and 40 points for third childbirth, which is considerable.

Second, when labor supply decisions in $t+1$ are allowed to be contingent not only on recent wages, but also on education, past full-time status and firm composition, we still find substantial heterogeneity along the recent wages distribution. Indeed, within groups of women with similar education and full-time status in $t-1$, we still find that those at the bottom of the wages distribution are 15 log-points (resp. 26 log-points and 32 log-points) less likely to have employment one year after the arrival of first (resp. second and third) child than those that ranked at the top of the distribution before chilbirth. Hence differences in education and full-time status, and differential sorting seem to account for a quite low share of the heterogeneity in labor supply response to childbirth at the extensive margin. While this contribution is still non-negligible (about a fifth of our initial estimate), we believe it tends to rule out the idea that mothers with low hourly wages being more likely to interrupt their careers is primarily the result of their unobserved preferences, rather than a decision based on the (relatively low) value of one hour spent in employment relative to one hour spent in home-production.

5.3 Motherhood penalties and fatherhood premias

We now turn to the consequences of childbirth on hourly wages. In order to do so, we focus on hourly wages 5 year after childbirth. Namely, we rely on our difference-in-difference approach to compare hourly wages growth between $t-1$ and $t+5$ of individuals who experienced childbirth during year t with that of invividuals of the adequate control group. Our choice to focus on $t+5$ stems from hourly wages at time t including (some part of) maternity leave allowances, which creates patterns in the data that we do not want to interpret as the effect of the arrival of a child on women's wage. We rely on OLS estimates of our difference-in-difference approach interacted with ranks in the recent wages distribution (see Subsection 4.3.2). Figures 4 and 5 display our estimates for women and men.

We find the arrival of the first child to have negative and significant impact on women's hourly wages 5 years after its birth, with the exception of women's that belong to both ends of the recent wages distribution. Our estimates suggest the effect of first childbirth on women's wages is about -5 log-points for the largest part of the recents wages distribution (Model 1). While the difference in the effect between percentiles of the distribution is not significant, our results lead to believe that the consequences of childbirth on mothers' wages might be slightly less harsh for both low and high earning women. Adding more controls for horizontal segregation – namely occupation, industry and firm composition – does not change much our estimates (Model 2). However, controling for experience, mobility and career interruptions lowers the effect (Model 3), which indicates, consistent with the rationale, that post-birth labor supply decisions of mothers, leading for instance to less human capital accumulation, might be a key driver of motherhood penalties. We also find that the arrival of a second child does not lead to statistically significant motherhood penalties (with respect to mothers of one), even though our confidence intervals are large so that we cannot reject the hypothesis that they generate economically significant wages losses.

When it comes to men, we find fathers that belong to the upper half of the

Figure 3 – Heterogeneity in the probability to remain in employment one year after childbirth: hourly wages only vs. hourly wages, education, full-time status and firm composition

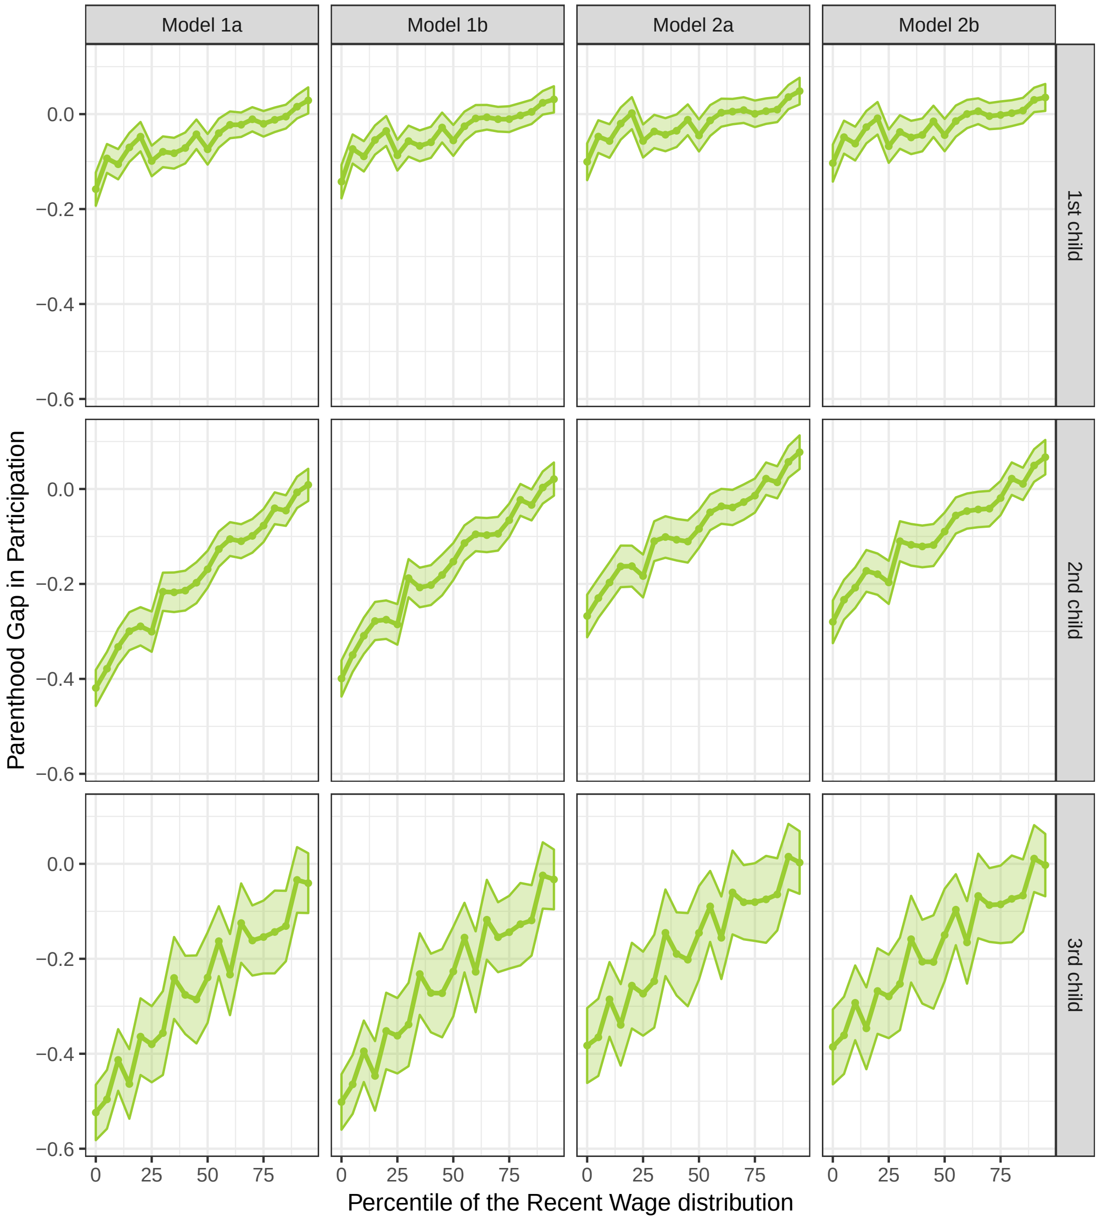

Estimates of the coefficients related to childbirth for women in a linear probability model that interacts a double-difference setting with gender and rank in the recent wages distribution (10). Outcome is a dummy for participating in the labor market at time $t + 1$. In Model 1, the difference-in-difference is only interacted with the recent wages distribution; in Model 2 it is also interacted with education and full-time status prior to childbirth. Models 1a and 2a include no controls; models 1b and 2b control for year, age, industry and 1-digit occupation within each gender *times* recent wages cell. Standard errors are clustered at the individual level. Sample includes individuals up to age 55 at time t.

recent wages distribution to experience faster hourly wages growth after the arrival of their first child than men that who do not have children (Model 1). This fatherhood premium might be as large as 7.5 log-points at the very top of the distribution. While the difference between our estimates is not significant, controlling for horizontal segregation (Model 2), and for experience accumulation and job mobility (Model 3) does lower them, which once again suggests that faster human capital accumulation due to slightly increased labor supply of fathers around childbirth may be at play here. The arrival of a second child does not generate significant fatherhood premias, with the exception of the very top of the recent wages distribution, but our confidence intervals are large so that we cannot reject economically significant effects.

With this evidence in mind, we turn to the consequences of childbirth on the gender pay gap, estimated by related triple-difference. Figure 6 displays our estimates. Consistent with previous results, the arrival of the first child has a significant negative impact on the gender pay gap – i.e. it widens the gap –, for all workers but those of the lowest end of the recent wages distribution. While variation along the distribution is not significant, our estimates suggest first child leads to larger gender gap among high achieving workers, up to 10 log-points 5 year after its birth, than among those that earn lower hourly wages, for which the effect amounts to 5 log-points. Controling for horizontal segregation (Model 2) does not change much our estimates, while accounting for experience and job mobility tends to lower them, which indicates once again that childbirth-related labor supply decisions might play a substantial role. We do not find significant effects for the arrival of a second child, but here again our confidence intervals are wide and cannot reject economically significant effects.

We also provide some evidence that the effect of some of our covariates may vary dramatically along the wages distribution. Namely, Figure 7 displays our estimates of the coefficients related to experience, career interruptions, job mobility and firm composition in Model 3. Experience is measured as the sum of hours of work between t and $t+5$, divided by the median duration of work for individuals employed full-time one year without any interruption (1820 hours); career interruptions are proxied by a dummy for spending at least one year between t and $t+5$ outside employment in the private sector. Job mobility is measured by a dummy for having different main employers[3] at time $t-1$ and $t+5$. Firm composition is measured by the share of part-time working women among employees of the same firme as i at time $t-1$.

Hourly wages growth is much more positively (resp. negatively) correlated with experience (resp. career interruptions) among high achieving workers than it is for their lower earning counterparts. This would be consistent with Dustmann and Meghir (2005) who find that returns to experience are larger among skilled workers than they are for unskilled ones. However, our coefficient cannot be taken neither as a correct estimate of returns to experience nor the causal effect of career interruptions on subsequent wages, because for forward-looking individuals, experience and past career interruptions reflect past labor supply decisions that were made based on expected future wages. We believe they do nevertheless bear some information. Firstly, this vast heterogeneity allows us to rationalize why low earning women, while being much more likely to reduce their labor supply at the arrival of a child, do not seem to encounter larger hourly wages penalties than their high earning counterparts. Secondly, it adds to our argument that heterogeneity in the opportunity cost of career interruptions is key to the understanding of mothers labor supply decisions. Indeed, if mothers are forward-looking, they would base their labor supply decision not only on their current hourly wages, but also on their expected future wages, which would be much more contingent on their current labor supply decisions for high achieving women than for their lower achieving counterparts. Thus, the difference between the value of one hour spent in the labor workforce and one hour spent outside employment should be much larger for those who earn high hourly wages, not only due to the direct contribution of their hourly wages, but also because spending time outside the labor market generates much larger wages losses when they get back to work.

Moving from one firm to the other coincides with hourly wages drops among workers that belong to the lowest part of the recent wages distribution, and with positive wages growth among those that previously earned high hourly wages. This would be the case if among low achieving workers, job mobility mostly results from job displacement, while high achieving workers use job mobility to benefit from competition between employers (Bagger et al., 2014). Lastly, we find workers that

[3] An individual main employer for a given year is the firm that pays him the highest labor earnings for that year.

belong to firms that employ high shares of part-time working women to have slower hourly wages growth than those that work in firms with lower shares of part-time working women. This could indicate that the sorting dimension investigated by Card, Cardoso, and Kline (2016); Coudin, Maillard, and To (2018) affects hourly wages not only in terms of levels, but also in terms of career progressions.

Figure 4 – Medium run consequences of childbirth on hourly wages: women

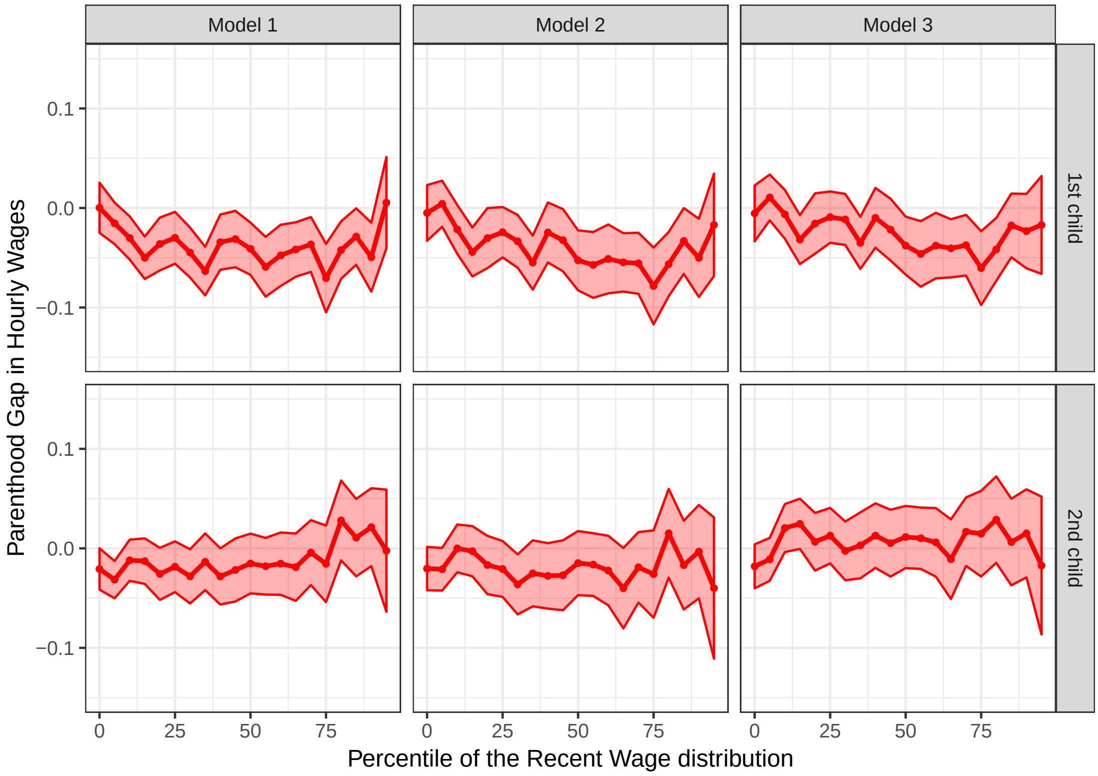

Estimates of the coefficients related to childbirth for women in hourly wages growth model that interacts a double-difference setting with gender and rank in the recent wages distribution (10). Outcome is a (log) hourly wages growth between $t-1$ and $t+5$. Model 1 includes not controls. Model 2 controls for year, age, industry, firm composition (share of part-time working women) and 1-digit occupation within each gender × recent wages cell. Model 3 includes all these controls plus experience between t and $t+5$, a dummy for having spent at least one year outside private sector employment, and having changed firm between $t-1$ and $t+5$. Standard errors are clustered at the individual level. Sample includes individuals up to age 55 at time t.

5.4 Children and sticky floors

Two distinct channels might be at play in this economically significant impact of childbirth on the gender pay gap. The first channel would correspond to men and

Figure 5 – Medium run consequences of childbirth on hourly wages: men

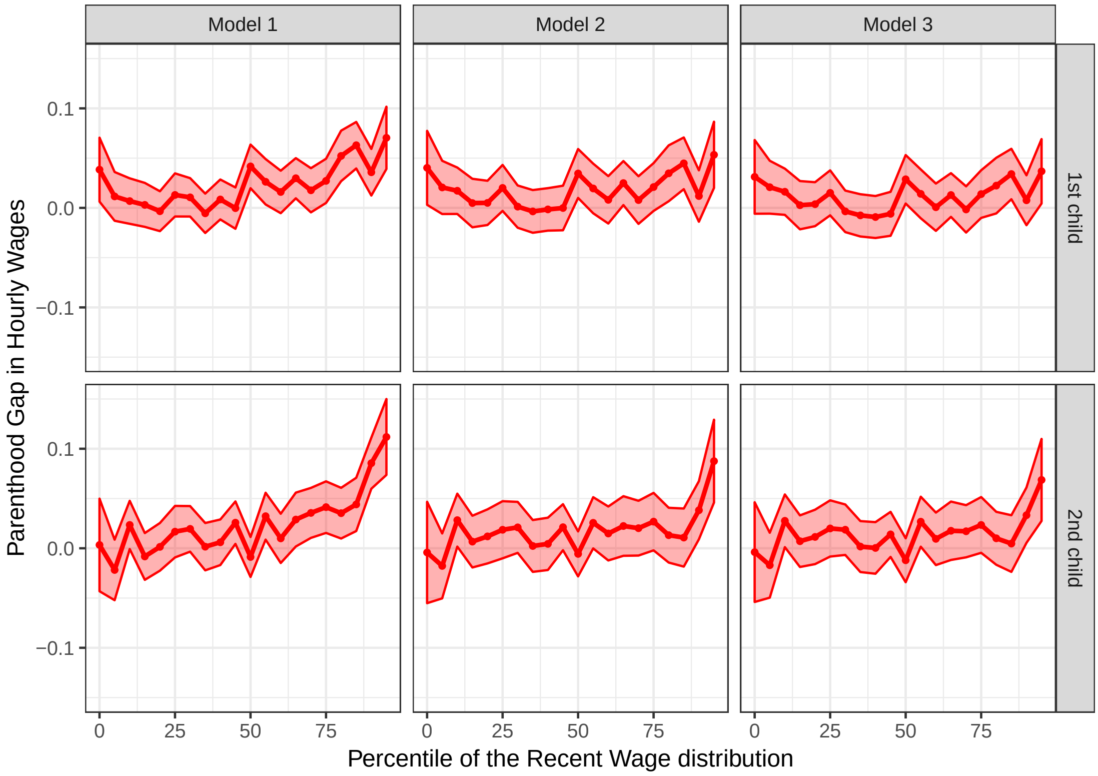

Estimates of the coefficients related to childbirth for men in hourly wages growth model that interacts a double-difference setting with gender and rank in the recent wages distribution (10). Outcome is a (log) hourly wages growth between $t-1$ and $t+5$. Model 1 includes not controls. Model 2 controls for year, age, industry, firm composition (share of part-time working women) and 1-digit occupation within each gender × recent wages cell. Model 3 includes all these controls plus experience between t and $t+5$, a dummy for having spent at least one year outside private sector employment, and having changed firm between $t-1$ and $t+5$. Standard errors are clustered at the individual level. Sample includes individuals up to age 55 at time t.

Figure 6 – Medium run consequences of childbirth on hourly wages: gender gap

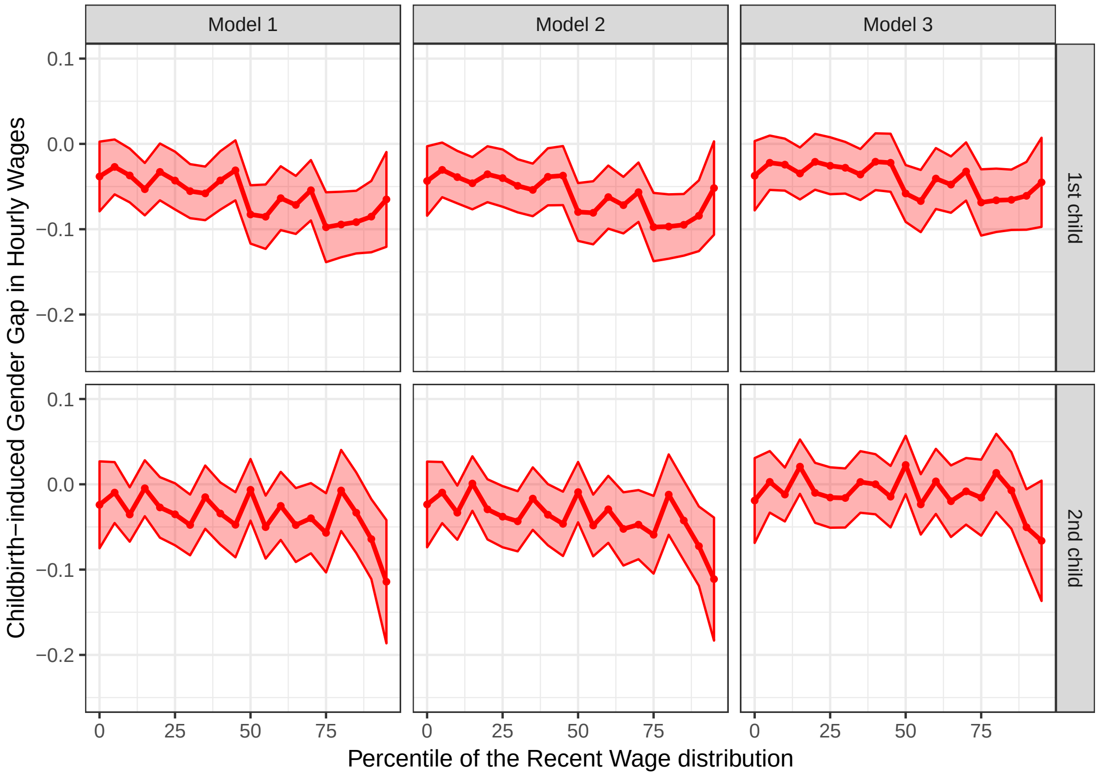

Estimates of the coefficients related to childbirth for women (men taken as a reference) in hourly wages growth model that interacts a double-difference setting with gender and rank in the recent wages distribution (11). Outcome is a (log) hourly wages growth between $t-1$ and $t+5$. Model 1 includes not controls. Model 2 controls for year, age, industry, firm composition (share of part-time working women) and 1-digit occupation within each recent wages cell. Model 3 includes all these controls plus experience between t and $t+5$, a dummy for having spent at least one year outside private sector employment, and having changed firm between $t-1$ and $t+5$. Standard errors are clustered at the individual level. Sample includes individuals up to age 55 at time t.

Figure 7 – Heterogeneity in returns to experience, career interruptions, firm composition and between-firm mobility

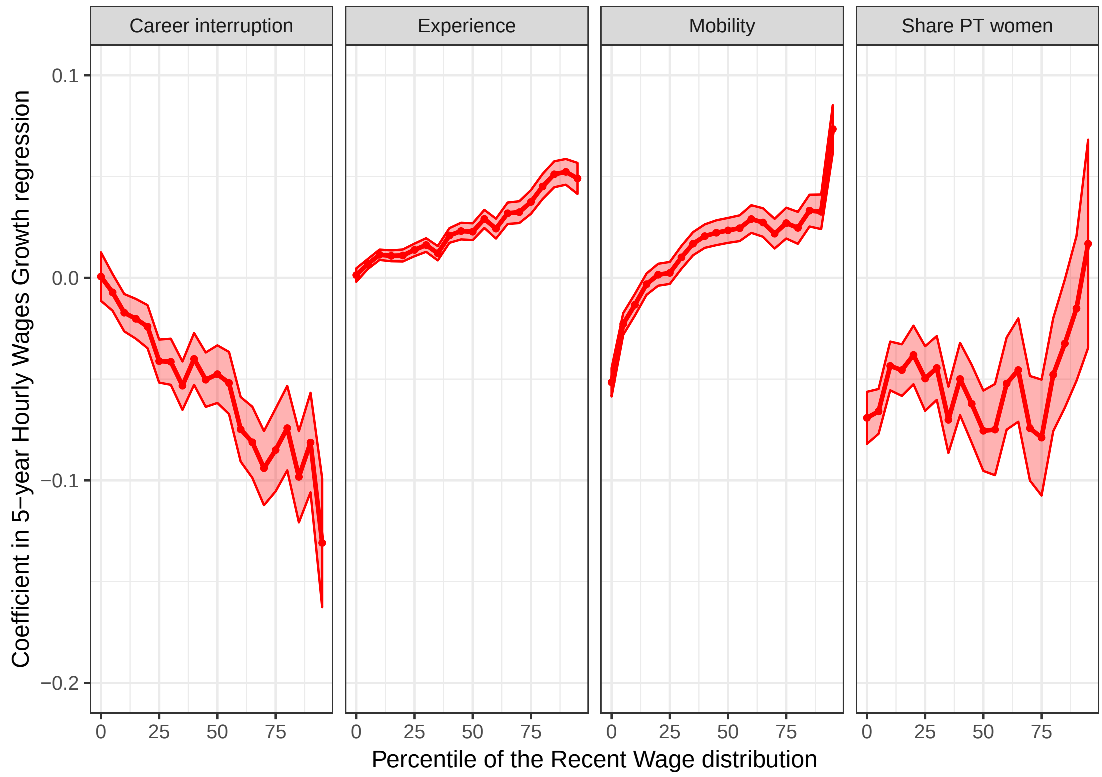

Estimates of the coefficients related to experience and career interruptions in hourly wages growth model that interacts a double-difference setting with gender and rank in the recent wages distribution (11). Outcome is a (log) hourly wages growth between $t-1$ and $t+5$. Model controls for year, age, industry, firm composition (share of part-time working women), 1-digit occupation, experience between t and $t+5$, a dummy for having spent at least one year outside private sector employment, and having changed firm between $t-1$ and $t+5$. Standard errors are clustered at the individual level. Sample includes individuals up to age 55 at time t.

women having diverging hourly wages levels due to the arrival of children. In this setting, childbirth would generate a permanent hourly wages shift between men and women when a child arrives, but afterwards they would experience similar career progressions, i.e. have parallel wages growth. The second channel would correspond to men and women having diverging slopes due to the arrival of a child. In this setting, childbirth may generate a instantaneous shift in women's hourly wages, but on top of that, because women spend less time on the labor market that do men after they have children, they would acquire less labor-market valued skills and thus experience slower wages growth than their male counterparts. Our previous results already suggest that the arrival of a child does generate a short-run shift in hourly wages for high achieving mothers, after which a catch-up may occur but is not sufficient for women to recover with their formerly similar male counterparts (Figures 1 and 2). The question is therefore whether in the long-run, several years after having a child, their hourly wages evolve as the same pace as those of men.

We investigate this issue by considering how the gender gap in hourly wages growth among individuals with children aged more than 6 at time t, that never experience new childbirth, differs from the gender gap in hourly wages growth among individuals that never have children. After age 6, most children attend school so that they should generate less time-constraints on there parents. Our approach can once again be thought of like a difference-in-difference, where our control group would be non-parents, except that we replace time dimension that is usual in difference-in-difference by a gender dimension. Figure 8 displays our OLS estimates, where outcome is 1-year hourly wages growth.

We find that in the lowest part of the recent wages distribution, the gender gap in hourly wages growth is larger among parents of children all aged more than 6 than among non-parents, while the difference is not significant among individuals placed in the highest half of the distribution (Model 1). This suggests that even quite long after they are born, and at ages where they are taken care of by schooling institutions, children would still affect their mothers' career progression. While the effect is not very large, about 1 log-point, it is worth noting that, as time goes by, those differences may cumulate and thus create very substantial gaps in hourly wages levels. Controling for horizontal segregation somehow lovers the effect

(Model 2), even though the difference is not statistically significant. This would be consistent with the view that the arrival of a child generating gender differentials in sorting across firms is a credible channel for this gap (Coudin, Maillard, and To, 2018). Surprisingly, further controling for experience and job mobility actually widens the gap (Model 3), which seem difficult to rationalize.

Overall, our findings may indicate that while among low earnings workers, childbirth does not seem to generate a very large instantaneous shift in hourly wages, which might be due to the binding constraint of the minimum wage, in the long-run children generate a sticky floor pattern, where mothers that already earn low hourly wages experience slower wages growth than their male counterparts, and are thus likely to remain stuck at the bottom of the wages distribution. Contrastingly, high achieving mothers may well experience instantaneous shifts in hourly wages, but several years after their last child is born the difference between their career progession and that of similar men is not larger than it is among non-parents, though this is not sufficient for them to catch-up with men.

5.5 Ceilings and floors among nonparents

Lastly, we show that even when they do not have children, men and women still have different career progressions. We show this by focusing on individuals that never have children in the data (they may still experience childbirth after 2015). We estimate gender differences in hourly wages growth all along the recents wages distribution, and display our estimates in Figure 9.

We find that at both ends of the distribution, women experience slower hourly wages growth than their male counterparts, while among median workers gender differences are not statistically significant (Model 1). While the difference is not necesarily very large, up to 1 log-point, once again if those gaps cumulate over time they can generate very substantial wages differentials. However, when controling for horizontal segregation (Model 2) and for experience and job mobility (Model 3), the difference among high hourly wages is no longer significant, even though confidence intervals are large enough to allow economically substantial differences. Differences among low wages earners remain barely significant even in the full control specification, which could point to a sticky floor effect that is not merely the consequence of children, and leads women with already quite low hourly wages

Figure 8 – Children-induced gender differences in hourly wages growth among parents of older children

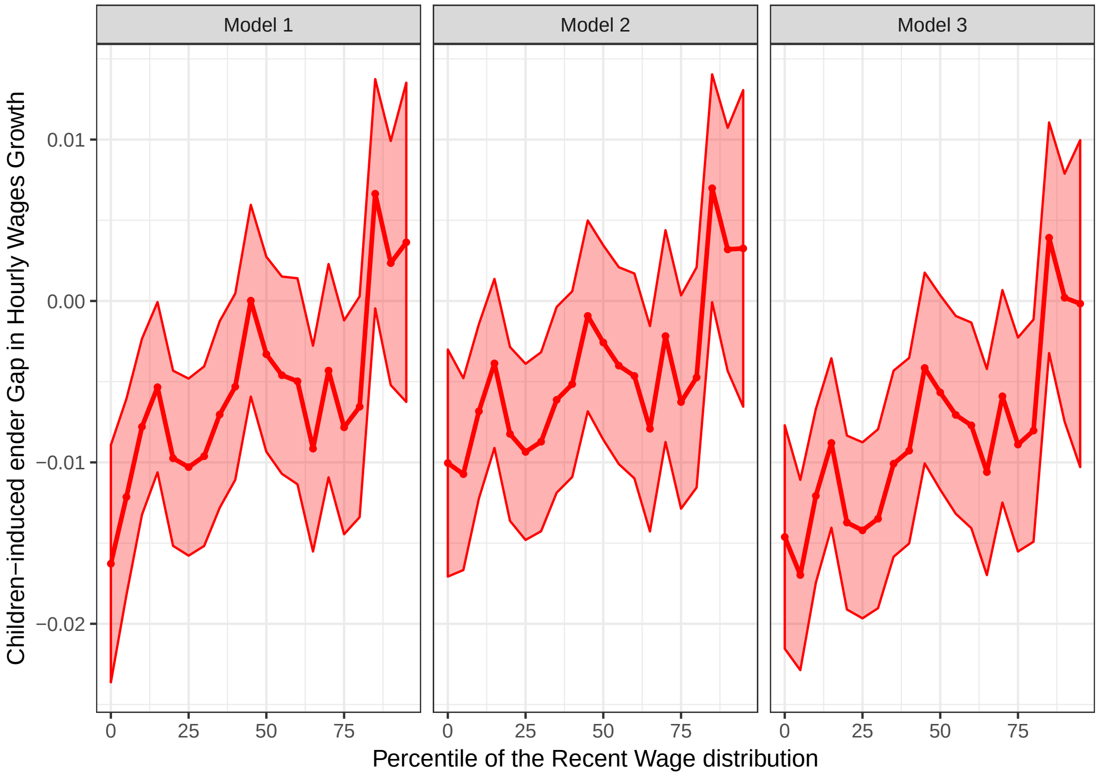

Estimates of the coefficients related to gender (female dummy) × having all children aged more than 6, relative to gender × never having children, interacted with location in the recent wages distribution, in a hourly wages growth model. Outcome is a (log) hourly wages growth between $t-1$ and t. Model 1 includes not controls. Model 2 controls for year, age, industry, firm composition (share of part-time working women) and 1-digit occupation within recent wages cell. Model 3 includes all these controls plus experience between t and t and having changed firm between $t-1$ and t. Standard errors are clustered at the individual level. Sample includes individuals up to age 60 at time t.

to be progressively left apart in the bottom of the distribution by formerly similar men. Note that while in the full specification, gender differences among top-earners are not statistically significant, the magnitude of our estimate is similar to that of low wages earners; this leaves room to a glass ceiling effect by which high achieving women experience less favorable career progression than their male counterparts, thus leaving the top of the distribution at a higher rate than them (Guvenen, Kaplan, and Song, 2014). Hence while the consequences of fertility may certainly play a part in the vertical segregation pointed out by Fortin, Bell, and Böhm (2017), it may however not be sufficient to fully account for it.

Figure 9 – Gender differences in hourly wages growth among nonparents

Estimates of the coefficients related gender (female dummy) hourly wages growth model that interacts gender and rank in the recent wages distribution. Outcome is a (log) hourly wages growth between $t-1$ and t. Model 1 includes not controls. Model 2 controls for year, age, industry, firm composition (share of part-time working women) and 1-digit occupation within recent wages cell. Model 3 includes all these controls plus experience between t and t and having changed firm between $t-1$ and t. Standard errors are clustered at the individual level. Sample includes individuals up to age 60 at time t.

6 Conclusion

This investigation of gender differences in career progressions pays a special attention to the effect of children on their parents' labor outcomes and to heterogeneity along the wages distribution, both of which have been pointed as key issues by recent contributions to the gender gap literature. Consistent with Kleven, Landais, and Søgaard (2018), we find the arrival of a child to have a very large negative impact on their mothers' labor earnings, and to coincides with slightly faster labor earnings growth of high earning fathers. However, in the short to medium run on which we focus, we show that this effect is primarily the result of labor supply decisions, and not so much of hourly wages. Furthermore, this effect is very heterogeneous, because mothers that previously earned low hourly wages are much more likely to leave the labor market or to reduce hours of work than those that earned high wages. Further analysis suggests that while intrinsic preferences towards family and career are certainly at play, this pattern is likely to be driven by contemporaneous incentives that affect the opportunity cost of carreer interruptions.

Contrastingly, the effect of childbirth on their mothers' hourly wages is quite homogeneous, or very slightly larger for high achieving mothers; fathers who previously earned high hourly wages may also experience a slight fatherhood premium. As a result, in the medium run, the gender gap in hourly wages widens more among high wages earners than among low wages earners: children are therefore likely to contribute to a substantial amount to the women's underrepresentation at the top of the distribution pointed out by Fortin, Bell, and Böhm (2017). This may seem puzzling given that career interruptions and labor supply reductions among mothers are less frequent for highly paid women than for poorly paid ones. We reconcile these findings by showing that returns to experience, and hourly wages losses due to career interruptions are presumably larger at the top of the wages distribution than at its bottom, which both explains why high achieving mothers would be more reluctant to spend time outside the labor workforce and why them spending less time outside employment still has more consequences in terms of their hourly wages. This would additionally be consistent with the view that highly paid occupation have more non-linear pay structures (Goldin, 2014), so that family-work conciliation is more difficult among them and leads to larger motherhood penalties (Bütikofer, Jensen, and Salvanes, 2018).

In the long run however, children do not seem to affect career progressions of high achieving mothers, but this is not sufficient for them to recover from the short to medium run negative consequences of childbirth. Reversely, low achieving mothers have slower hourly wages growth long after the arrival of their last child, while they presumably experienced lower negative consequences on their hourly wages levels in the short to medium run. This may generate a children-related sticky floor pattern, where low wages earning mothers tend to remain stuck at the bottom of the wages distribution. Lastly, we show that even among non-parents, men and women differ in their career progressions at both ends of the wages distribution, which suggests that dynamic vertical segregation (Fortin, Bell, and Böhm, 2017), i.e. glass ceilings and sticky floors is not merely the result of children-related constraints and decisions.

A challenging task would be to reconcile these findings in a full decomposition of the gender pay gap, that would estimate jointly the contribution of children and the contribution of vertical segregation. This would probably rely on a full structural specification of a lifecycle model with endogeneous fertility decisions and heterogeneous abilities and preferences (Adda, Dustmann, and Stevens, 2017) that is beyond the scope of this paper, so that we leave it for further research.

Children-related labor supply decisions of mother being seemingly driven by contemporaneous incentives suggests that changing the environment in which such decisions are made can have first-order consequences on career progressions of mother, and thus on the gender pay gap. These changes can be achieved for instance through paternity leaves, or cautious design of parental leave allowances (Piketty, 2005; Lequien, 2012; Joseph et al., 2013) or childcare subsidies (Givord and Marbot, 2015). Specifically, our results could indicate that parental leave allowances not being contingent on potential hourly wages is likely to pull low wages earning mothers out of the labor workforce, and therefore to lead to them having lower wages growth in the future, i.e. to contribute to a sticky floor effect. When it comes to high achieving women however, policy implications of our results do no seem obvious: the arrival of children does indeed contribute to underrepresentation of women among top-earners, but we find that highly paid mothers are very unlikely to leave employment, and do not substantially reduce their working hours after childbirth, so that it does not make sense to increase incentives for them to remain on the labor market.

Another consequence of this result is that being the result of past labor supply decisions, and specifically of children-related labor supply decisions, that are contingent on hourly wages, the inclusion of experience in cross-sectional Oaxaca (1973)-Blinder (1973) decompositions of the gender pay gap is likely to generate biased estimates. Precisely, it is likely to underestimate the unexplained part that is interpreted as discrimination. While researchers have already aknowledged this (Kunze, 2008), leading some of them to rely on exogeneous policy shocks to estimate to consequences of mothers' career interruptions (Lequien, 2012), we believe recalling this can be useful to policy makers when implementing public policies explicitly based on this econometric framework (Vaccaro, 2017).

References

Adda, J., C. Dustmann, and K. Stevens. 2017. "The Career Costs of Children." *Journal of Political Economy* 125:293–337.

Afsa, C., and S. Buffeteau. 2006. "L'activité féminine en France : quelles évolutions récentes, quelles tendances pour l'avenir ?" *Économie et Statistique* 398-399:85–97.

Albrecht, J., A. Björklund, and S. Vroman. 2003. "Is There a Glass Ceiling in Sweden?" *Journal of Labor Economics* 21:145–177.

Anderson, D.J., M. Binder, and K. Krause. 2003. "The Motherhood Wage Penalty Revisited: Experience, Heterogeneity, Work Effort, and Work-Schedule Flexibility." *Industrial and Labor Relations Review* 56:273–294.

Angrist, J.D., and W.N. Evans. 1998. "Children and Their Parents' Labor Supply: Evidence from Exogenous Variation in Family Size." *The American Economic Review* 88:450–477.

Arulampalam, W., A.L. Booth, and M.L. Bryan. 2007. "Is There a Glass Ceiling over Europe? Exploring the Gender Pay Gap across the Wage Distribution." *Industrial and Labor Relations Review* 60:163–186.

Bagger, J., F. Fontaine, F. Postel-Vinay, and J.M. Robin. 2014. "Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics." *American Economic Review* 104:1551–1596.

Becker, G. 1981. *A Treatise on the Family*. Cambridge: Harvard University Press.

Bertrand, M., E. Duflo, and S. Mullainathan. 2004. "How Much Should We Trust Differences-In-Differences Estimates?" *The Quarterly Journal of Economics* 119:249–275.

Bertrand, M., C. Goldin, and L.F. Katz. 2010. "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors." *American Economic Journal: Applied Economics* 2:228–55.

Blinder, A.S. 1973. "Wage Discrimination: Reduced Form and Structural Estimates." *The Journal of Human Resources* 8:436–455.

Bütikofer, A., S. Jensen, and K.G. Salvanes. 2018. "The role of parenthood on the gender gap among top earners." *European Economic Review*, pp. .

Carbonnier, C. 2014. "The influence of taxes on employment of married women, evidence from the French joint income tax system." Sciences Po publications No. 23, Sciences Po.

Card, D., A.R. Cardoso, and P. Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women *." *The Quarterly Journal of Economics* 131:633–686.

Charnoz, P., E. Coudin, and M. Gaini. 2011. "Changes in the French Wage Distribution 1976-2004: Inequalities within and between Education and Experience Groups." Working paper INSEE.

Choné, P., D. Le Blanc, and I. Robert-Bobée. 2004. "Offre de travail féminine et garde des jeunes enfants." *Économie & prévision* 162:23–50.

Christofides, L.N., A. Polycarpou, and K. Vrachimis. 2013. "Gender wage gaps, 'sticky floors' and 'glass ceilings' in Europe." *Labour Economics* 21:86 – 102.

Coudin, E., S. Maillard, and M. To. 2018. "Family, firms and the gender wage gap in France." IFS Working Papers No. W18/01, Institute for Fiscal Studies, Jan.

de la Rica, S., J.J. Dolado, and V. Llorens. 2008. "Ceilings or floors? Gender wage gaps by education in Spain." *Journal of Population Economics* 21:751–776.

Deaton, A.S. 1997. "Econometric Issues for Survey Data." In T. W. Bank, ed. *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. The Johns Hopkins University Press.

Deaton, A.S., and C. Paxson. 1994. "Saving, Growth, and Aging in Taiwan." In *Studies in the Economics of Aging*. National Bureau of Economic Research, Inc, NBER Chapters, pp. 331–362.

Dustmann, C., and C. Meghir. 2005. "Wages, Experience and Seniority." *The Review of Economic Studies* 72:77–108.

Etienne, J.M., and M. Narcy. 2010. "Gender Wage Differentials in the French Nonprofit and For-profit Sectors: Evidence from Quantile Regression." *Annals of Economics and Statistics*, pp. 67–90.

Firpo, S., N.M. Fortin, and T. Lemieux. 2009. "Unconditional Quantile Regressions." *Econometrica* 77:953–973.

Fortin, N.M., B. Bell, and M. Böhm. 2017. "Top earnings inequality and the gender pay gap: Canada, Sweden, and the United Kingdom." *Labour Economics* 47:107 – 123, EALE conference issue 2016.

Givord, P., and C. Marbot. 2015. "Does the cost of child care affect female labor market participation? An evaluation of a French reform of childcare subsidies." *Labour Economics* 36:99 – 111.

Gobillon, L., D. Meurs, and S. Roux. 2015. "Estimating Gender Differences in Access to Jobs." *Journal of Labor Economics* 33:317–363.

Goldin, C. 2014. "A Grand Gender Convergence: Its Last Chapter." *American Economic Review* 104:1091–1119.

Guvenen, F., G. Kaplan, and J. Song. 2014. "The Glass Ceiling and The Paper Floor: Gender Differences among Top Earners, 1981-2012." NBER Working Papers No. 20560, National Bureau of Economic Research, Inc.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song. 2017. "Heterogeneous Scarring Effects of Full-Year Nonemployment." *American Economic Review* 107:369–73.

—. 2016. "What do data on millions of US workers reveal about life-cycle earnings risk?" mimeo.

Huss, M.M. 1990. "Pronatalism in the Inter-War Period in France." *Journal of Contemporary History* 25:39–68.

Joseph, O., A. Pailhé, I. Recotillet, and A. Solaz. 2013. "The economic impact of taking short parental leave: Evaluation of a French reform." *Labour Economics* 25:63 – 75, European Association of Labour Economists 24th Annual Conference, Bonn, Germany, 20-22 September 2012.

Juhn, C., and K. McCue. 2017. "Specialization Then and Now: Marriage, Children, and the Gender Earnings Gap across Cohorts." *Journal of Economic Perspectives* 31(1):183–204.

Kleven, H., C. Landais, and J.E. Søgaard. 2018. "Children and Gender Inequality: Evidence from Denmark." NBER Working Papers No. 24219, National Bureau of Economic Research, Inc, Jan.

Koenker, R., and G. Bassett. 1978. "Regression Quantiles." *Econometrica* 46:33–50.

Korenman, S., and D. Neumark. 1992. "Marriage, Motherhood, and Wages." *The Journal of Human Resources* 27:233–255.

Kunze, A. 2008. "Gender wage gap studies: consistency and decomposition." *Empirical Economics* 35:63–76.

Lequien, L. 2012. "The Impact of Parental Leave Duration on Later Wages." *Annals of Economics and Statistics*, pp. 267–285.

Mason, K.O., W.M. Mason, H.H. Winsborough, and W.K. Poole. 1973. "Some Methodological Issues in Cohort Analysis of Archival Data." *American Sociological Review* 38:242–258.

Meurs, D., A. Pailhé, and S. Ponthieux. 2010. "Child-related Career Interruptions and the Gender Wage Gap in France." *Annals of Economics and Statistics*, pp. 15–46.

Oaxaca, R. 1973. "Male-Female Wage Differentials in Urban Labor Markets." *International Economic Review* 14:693–709.

Piketty, T. 2005. "L'impact de l'allocation parentale d'éducation sur l'activité féminine et la fécondité en France, 1982-2002." *Les Cahiers de l'INED*, pp. 79–109.

Powell, D. 2016. "Quantile regression with nonadditive fixed effects." mimeo.

Rodrigues, B., and V. Vergnat. 2016. "The impact on wages and worked hours of childbirth in France." Working Paper No. 2016-48, Bureau d'économie théorique et appliquée (BETA).

—. 2018. “The time and the transitions back to work in France after maternity.” Working Paper No. 2018-14, Bureau d’économie théorique et appliquée (BETA).

Rosental, P.A. 2010. “Politique familiale et natalité en France : un siècle de mutations d’une question sociétale.” *Santé, Société et Solidarité* 9:17–25.

Rosenzweig, M.R., and K.I. Wolpin. 1980. “Life-Cycle Labor Supply and Fertility: Causal Inferences from Household Models.” *Journal of Political Economy* 88:328–348.

Simonsen, M., and L. Skipper. 2006. “The costs of motherhood: an analysis using matching estimators.” *Journal of Applied Econometrics* 21:919–934.

Vaccaro, G. 2017. “Can Policy Directly Reduce the Unexplained Gender Wage Gap? Evidence from Switzerland.” mimeo.

Waldfogel, J. 1997. “The Effect of Children on Women’s Wages.” *American Sociological Review* 62:209–217.

—. 1995. “The Price of Motherhood: Family Status and Women’s Pay in Young British Cohort.” *Oxford Economic Papers* 47:584–610.

—. 1998. “Understanding the ”Family Gap” in Pay for Women with Children.” *The Journal of Economic Perspectives* 12:137–156.

Wilner, L. 2016. “Worker-firm matching and the parenthood pay gap: Evidence from linked employer-employee data.” *Journal of Population Economics* 29:991–1023.

Wood, R.G., M.E. Corcoran, and P.N. Courant. 1993. “Pay Differences among the Highly Paid: The Male-Female Earnings Gap in Lawyers’ Salaries.” *Journal of Labor Economics* 11:417–441.

A Childbirth imputation

We combine data issued from administrative birth records with census data in order to deal with the incompleteness of administrative birth records for individuals born October 2nd and 3rd in our dataset. Secifically, (part of) birth records are missing for these individuals between 1982 and 1997. Our strategy is to take information from the 1990 and 1999 censuses in order to fill the gap.

For each individual in our sample, our data provides us with:

- year of birth for 1st to 12th child that appear in birth records as of 1967;
- year of birth for 1st to 12th child as declared in the 1990 census;
- year of birth for 1st to 12th child as declared in the 1999 census.

Information from birth records being only available as of 1967 creates a left-censoring of our data. However, because we are mostly interested in individuals that give birth between 2005 and 2015, we do not try to deal with this issue. Our main goal is to fill the gap in administrative records between 1982 and 1997 for half of the sampled individuals, in order to ensure sufficient sample size for our analysis.

In order to do so, for each individual i of the problematic half of the sample, we first impute year of first childbirth according to the following approach:

- if the first childbirth in birth records occurs before 1982, we take it as first childbirth;
- else:
  - if the minimum of years of childbirth she declared in the 1990 census is superior or equal to 1982, we take as year of first childbirth the minimum of these years and the year of first childbirth as it appears in birth records;
  - else:
    * if the minimum of years of childbirth she declared in the 1999 census is superior or equal to 1982, we take as year of first childbirth the minimum of these years and the year of first childbirth as it appears in birth records;

    * else:
        - if she has children according to birth records we take as year of first childbirth the year of first childbirth in birth records;
        - else we consider her to be without child.

We then consider nth childbirth, with $n > 1$: we consider it to be the minimum of years of childbirth within both birth records and both censuses data, among years of birth that are superior to computed year of $n - 1$th childbirth.

Hence our rather simple approach does not take into account twin births, and more generally does not allow individuals to experience more than one childbirth a year. With this caveat in mind, we still show that for the sake of this paper, our approach matches quite well the historical pattern is the good part of the sample. Figure 10 plots the number of childbirths for each year since 1968, by rank of childbirth, for both parts of the sample, first relying only on data from birth records (left panel) and the relying on our approach (right panel). While we still slightly underestimate first childbirths that occur at the beginning of the 1980s or in the late 1990s in the problematic part of the sample, our approach does a reasonable job at matching the patterns that we observe in the clean part, especially over 2005-2015 that is the period on which our analysis focuses.

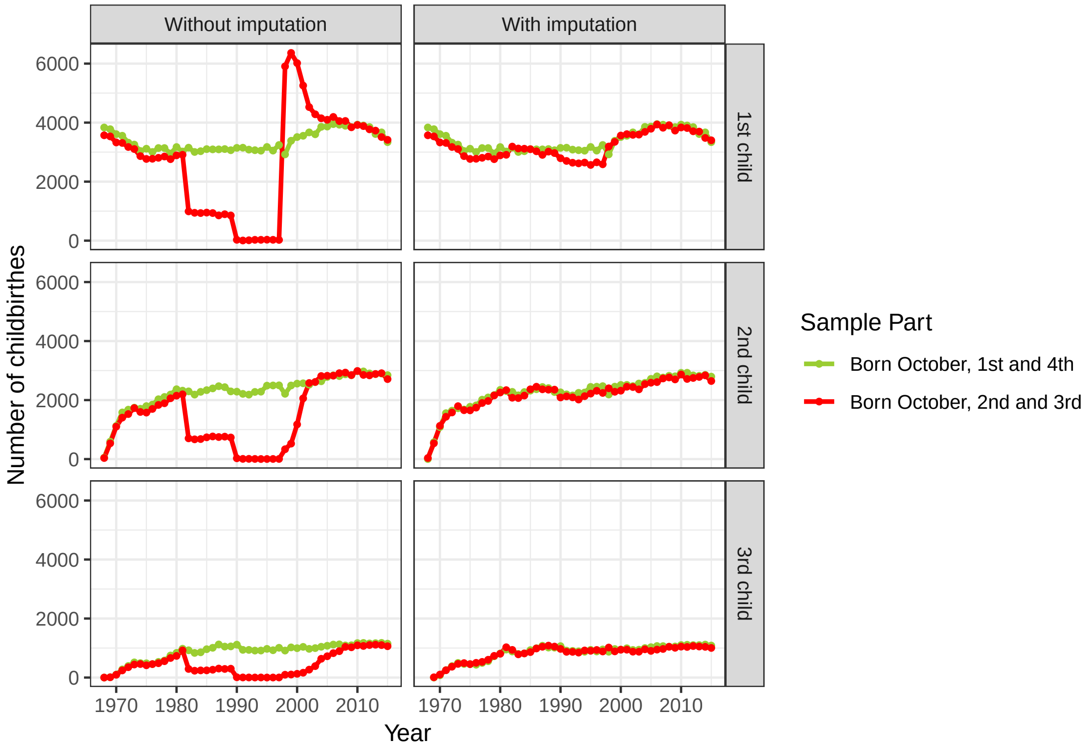

Figure 10 – Imputation of childbirths for individuals born October 2nd and 3rd

B Robustness checks: descriptive evidence

Figure 11 – Consequences of first childbirth on women's labor outcomes: by rank in the recent wages distribution. Median changes estimates.

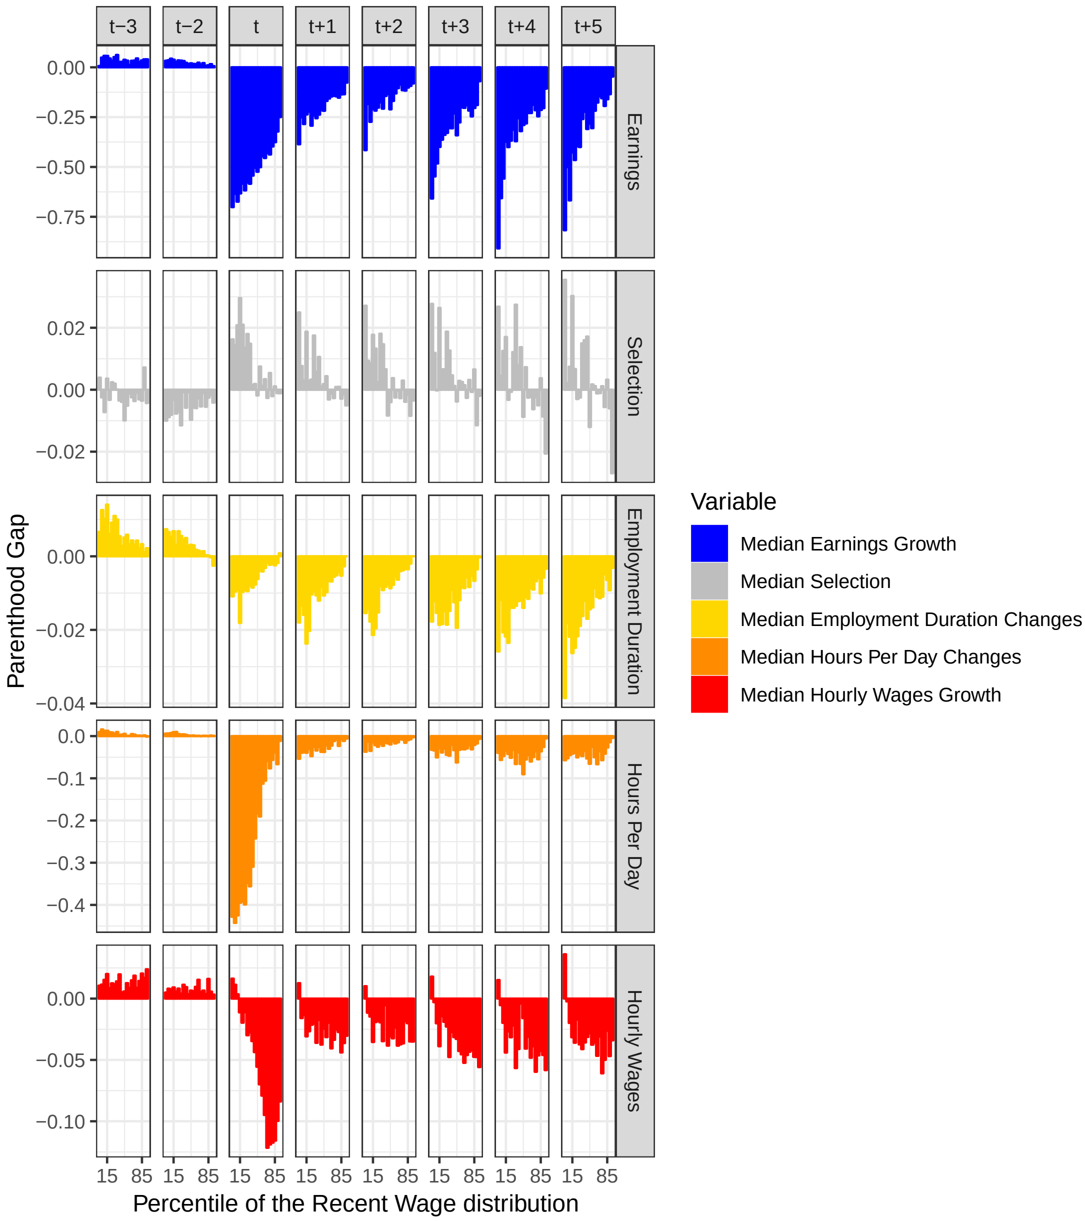

Figure 12 – Consequences of second childbirth on women's labor outcomes: by rank in the recent wages distribution. Median changes estimates.

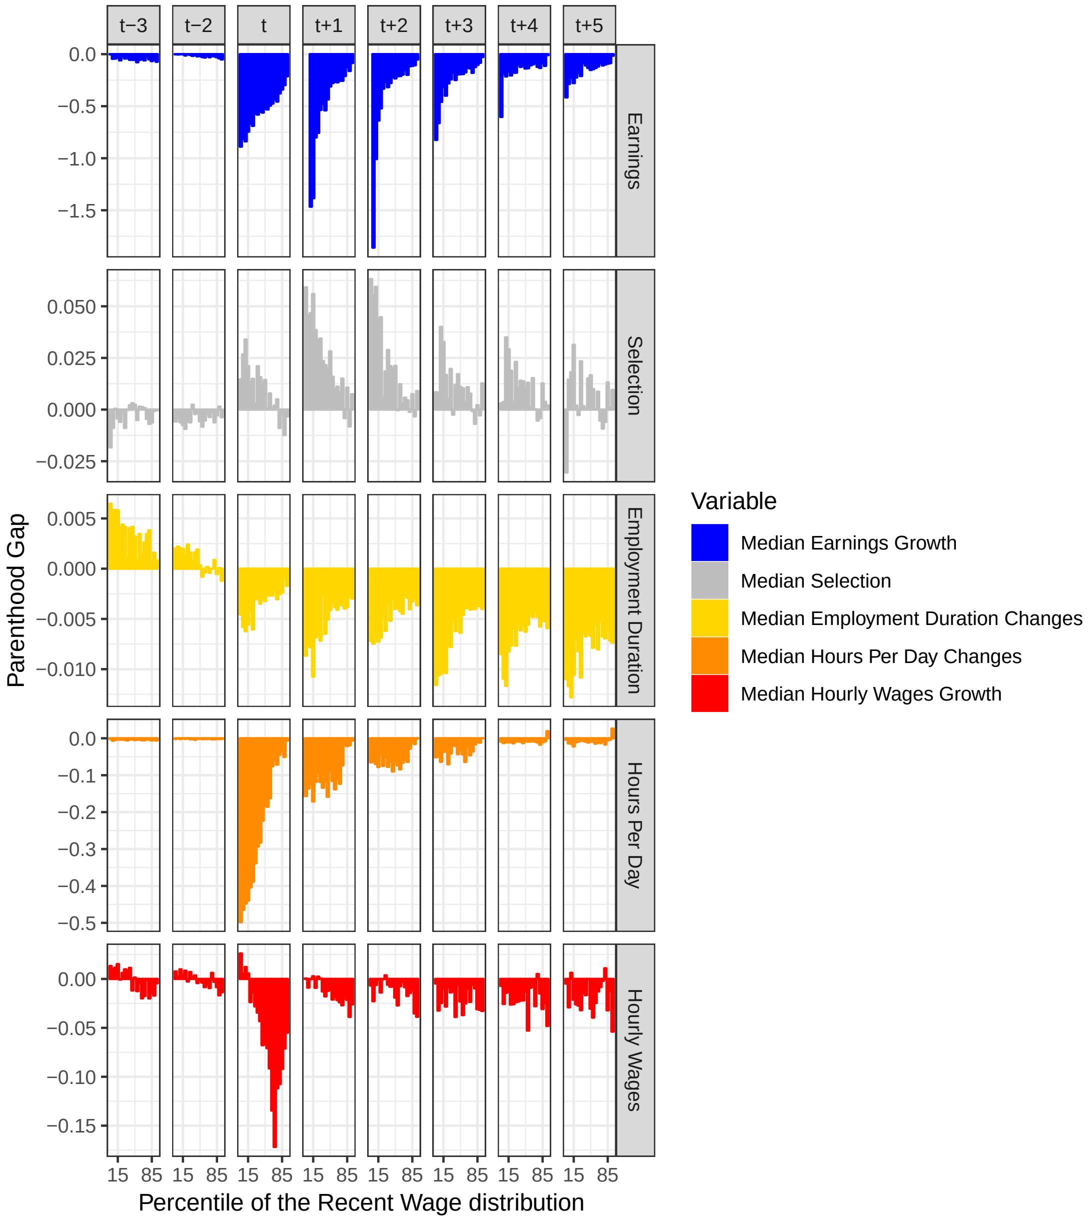

Figure 13 – Consequences of third childbirth on women's labor outcomes: by rank in the recent wages distribution. Median changes estimates.