
CEILINGS AND FLOORS. THE GENDER PAY GAP OVER THE LIFE CYCLE 2005-2012 ¹

Pierre PORA(*)(**), *Lionel WILNER*(*)(**)

(*) *Insee*

(**) *CREST*

pierre.pora@insee.fr

Mots-clés. Gender pay gap, glass ceilings, sticky floors.

Résumé

Nous revisitons des résultats récents qui mettent en évidence deux canaux qui contribuent à la persistance d'écarts de rémunération entre hommes et femmes : la sous-représentation des femmes au sommet de la distribution de salaire, et l'augmentation de l'écart au long du cycle de vie, liée à la parentalité. Nous employons une approche descriptive développée pour l'étude des dynamiques de revenu salarial, et l'appliquons à des données longitudinales issues de sources administratives. Les deux canaux sont relativement indépendants l'un de l'autre : les écarts de salaire entre hommes et femmes parmi les très hauts salaires sont bien plus importants que pour le reste des salariés, mais les salariés les mieux rémunérés ne diffèrent pas des autres pour ce qui est de l'écart en progression de salaire ou des pénalités liées aux naissances. Les écarts entre hommes et femmes parmi les très hauts salaires apparaissent dès l'entrée sur le marché du travail, peut-être du fait du coût anticipé des interruptions de carrière.

Abstract

We revisit recent research that has highlighted two channels that contribute massively to the aggregate gender pay gap: dramatic underrepresentation of women at the very top of the distribution, and increasing gender gap over the lifecycle, mainly due to childbirth. We rely on a descriptive framework developed for the study of earnings dynamics and apply it to a large French longitudinal dataset. Those two channels are seemingly unrelated: while there is a massive wage gap at the very top of the wage distribution, top earners do not depart from the rest of the workers when it comes to gender differences in wage growth and motherhood penalties over the lifecycle. Gender differences among top earners arise as individuals enter the labor market, possibly because of anticipated career interruptions costs.

¹We are grateful to Thierry Magnac for his suggestions that are at the origin of this project. We thank Nila Ceci-Renaud, Élise Coudin, Sébastien Roux and Grégory Verdugo for useful suggestions. All errors and opinions are ours.

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 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 Sogaard, 2018). While the evidence sustaining each claim is substantial, one is led to wonder whether those are related or not. Namely, does underrepresentation of women at the very top of the earnings distribution stem from childbirth or does it constitute a different stylised fact in itself ?

Answering this question is valuable to the design of efficient gender equality-oriented public policies, because it indicates whether specific events – like childbirth – or groups – like top earners – are to be more specifically targeted by the policy maker. Several measures intended to increase women representation at the top of the job ladder, like gender quotas in governing boards, have indeed proved disappointing (see Fortin, Bell, and Böhm, 2017 for a review). Getting to know if there is a limited stage of the lifecycle at which action is better intended can therefore be of significant help.

In this paper, we examine the relationship between motherhood penalties and gender differences among top earners, and more generally investigate how heterogeneous the gender gap may be along the wage distribution. First, we revisit evidence on the heterogeneity of the gender pay gap along the wage distribution and find that: (i) there exists a very large wage gap among top earners (namely the top 1% of a cohort at a given point in time); (ii) the gender wage gap is U-shaped, i.e. larger at both ends of the wage distribution ; (iii) these patterns are not likely explained by childbirth or career interruptions. In a second part, we focus on the distributions of individual wage changes and find suggestive evidence that (iv) childbirth explain most of the differences in wage growth between genders; (v) gender differences in wage growth among top earners do not depart from the rest of the distribution; (vi) career interruptions have far worse consequences for top earners than they do for the rest of the distribution. These results imply that the impressive gender gap among top earners is plausibly related to pre-labor market disparities or choices made at labor market entry, which in turn can reflect anticipations of the cost of time spent out of the labor market.

Dealing with these issues – top earnings inequality and differences that may occur at different points of the lifecycle – requires large and detailed longitudinal datasets that are not always 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. We focus on individuals working in the private sector from 2002 to 2015. Our approach builds on a framework recently developed for the study of labor earnings dynamics and risk (Güvenen, Ozkan, and Song, 2014; Güvenen et al., 2016, 2017), that allows for unrestricted heterogeneity in labor earnings trajectories along the earnings distribution.

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 presents our results, and section 6 concludes.

2 Literature review

Heterogeneity in the gender pay has been regarded as evidence of a *glass ceiling* (resp. *sticky floor*) when gender differences turn out to be larger at the top (resp. at the bottom) of the earnings and wage distributions. Relying on quantile regressions, Albrecht, Björklund, and Vroman (2003), Gardeazabal and Ugidos (2005), Etienne and Narcy (2010) find the gap to widen at the top of the distribution, while de la Rica, Dolado, and Llorens (2008) show it to widen at the bottom of the distribution among low-educated workers. Applying this quantile regression framework to several European countries, Arulampalam, Booth, and Bryan (2007) prove the widening of the gap at the top of the distribution to be quite general, whereas larger gaps at the bottom are only found in a few countries. Christofides, Polycarpou, and Vrachimis (2013) also estimate the gender wage gap at different quantiles of the distribution over several European countries. They find the pattern to be rather country-specific, so that in some countries women face a glass ceiling while in others they face a sticky floor (and in some they face both). While Arulampalam, Booth, and Bryan (2007) relates those between-countries differences to child-care provision, Christofides, Polycarpou, and Vrachimis (2013) link them to wage-setting institutions.

Interpretation of these quantile regression estimates in terms of glass ceiling and sticky floor is nevertheless questionable: if earnings and wages are not linear in rank, then a constant gender difference in access to jobs along the distribution is sufficient to generate those patterns, as pointed out by Gobillon, Meurs, and Roux (2015). However they do find that gender differences in access function are indeed increasing with rank in the job ladder, which would be consistent with a glass ceiling.

Because the wage and earnings distributions have fat tails, gender differences at the very top of the distribution may result in substantial aggregate gender pay gap. In a recent paper, Fortin, Bell, and Böhm (2017) show the huge underrepresentation of women among top earners (defined by their rank in the overall earnings distribution) to explain a substantial part of the aggregate gender gap. Furthermore, this underrepresentation is not accounted for by traditional factors such as education or occupational sorting. However this rests on cross-sectional evidence, which makes it difficult to know whether this huge underrepresentation stems from pre-labor market differences in unobserved characteristics, is determined once and for all at the beginning of a career, or is the consequence of gender differences in the probability of accessing and leaving each rank of the earnings distribution that vary over the lifecycle. 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 the probability of moving up in the distribution is higher for men than 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.

There is indeed evidence that the gender pay gap does not widen within cohort only for top earners. Indeed, Manning and Swaffield (2008), Bertrand, Goldin, and Katz (2010) and Goldin et al. (2017) find the gender pay gap to be negligible at labor market entry, and to widen substantially during the first years of a career. However evidence is not clear-cut since Morgan (1998), Kunze (2003), Kunze (2005), Weinberger and Kuhn (2010) and Weinberger (2011) for instance find the gender gap not to widen as cohorts grow of age. Furthermore Manning and Robinson (2004) show no differences in earnings growth between men and women with continuous employment.

Another strand of the literature has long investigated how men and women's labor market outcomes may diverge over the lifecycle, especially because childbirth 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 has in particular focused on the endogeneity of women’s fertility decisions, resorting 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), family background (Korenman and Neumark, 1992), siblings sex mix (Angrist and Evans, 1998) as instruments for children, propensity score matching (Simonsen and Skipper, 2006), two-way fixed-effects to eliminate the contribution of the sorting of individuals between firms (Wilner, 2016). However, using simple event-study techniques, Kleven, Landais, and SØgaard (2018) find the causal effect of childbirth, estimated thanks to sex-mix instruments, not to differ much from their initial estimate.

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. 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 wage structure, but Kleven, Landais, and SØgaard (2018) find very limited heterogeneity in terms of educational levels. Hence it remains difficult to assess how glass ceilings, sticky floors and paper floors are related to motherhood.

3 Data

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,² 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 employee 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 individual 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.

Individuals 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 incomplete during for some part of the sample during the 1990s, which has been documented (see Wilner, 2016). We only use information on childbirth that have occurred between 2002 and 2015. 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 metropoli-

²The absence of a DADS as well as incorrect or missing answers are punished with fines

tan 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 wage distribution. We impose in particular that these individuals are present at least to years between $t - 5$ and $t - 2$ on top of being present in $t - 1$. To deal with very low annual earnings, we focus on individuals earning more than $1/8$ of the annual minimum wage \underline{w}_t , as Guvenen et al. (2017) do. We also winsorize labor earnings at quantile of order 0.99999, in order to avoid issues related to potential outliers. In the end, our sample gathers over 6 million individuals-years observations, corresponding to more than 900 000 workers.

In Table 1, we give some descriptive statistics on the successive steps of the selection of "relatively stable" workers. First comes the censoring at $1/8\underline{w}_t$. 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

	Base sample				Censoring				Final sample			
	Women	Men	Women	Men	Women	Men	Women	Men	Women	Men	Women	Men
N_{it}	9 940 285	5 571 225	3 910 929	5 254 842	9 165 771	5 254 842	2 843 890	3 979 832	6 823 722	2 843 890	3 979 832	6 823 722
N_i	1 243 794	682 131	526 979	631 084	1 158 063	631 084	414 732	529 655	944 387	414 732	529 655	944 387
	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

4 Empirical analysis

Our empirical approach is largely derived from the descriptive framework developed by Guvenen et al. (2016, 2017) for the study of labor earnings dynamics and risk. In particular, we rely on non-parametric estimations of the distribution of future earnings and hourly wages condition on gender, age and rank in the overall distribution of past wages.

Let denote hourly wages for individual i on year $t = 1 \dots T$ by \tilde{w}_{it} . We consider a normalized version of earnings net of age effects. Hence, we regress individual log hourly wages on a set of age, (year) period and (year-of-birth) cohort dummies:

$$\ln(\tilde{w}_{it}) = \ln(w_0) + \sum_c \alpha_c \mathbb{1}_{\text{cohort}_i=c} + \sum_a \beta_a \mathbb{1}_{\text{age}_{it}=a} + \sum_j \gamma_j \mathbb{1}_{t=j} + \epsilon_{it} \quad (1)$$

The inclusion of year dummies is a slight difference with Guvenen et al. (2016, 2017). We introduce them to control for any disruption caused by methodological changes in the fabrication process of the DADS panel which occurred in 2002, 2009 and 2013.

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 α , β and γ 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 transformation³ 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 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 (α, β, γ) are identified provided that $\alpha_{\underline{c}} = 0$, $\beta_{\underline{a}} = 0$ and $\sum_{t=1}^T \gamma_t(t-1) = 0$. The corresponding transformation of time dummies d_j writes as follows:

$$d_t^* = d_t - [(t-1)d_2 - (t-2)d_1]$$

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

4.1 Distribution of recent wages

We aim at comparing workers with similar hourly wage histories. To do so, we introduce a measure of recent wages W_{it} similar to the one used by Guvenen et al. (2016, 2017) for earnings. This measure approximates average hourly wages between $t-5$ and $t-1$ net of age effects:

$$W_{it} = \frac{\sum_{\tau=t-5}^{t-1} \tilde{w}_{i\tau}}{\sum_{\tau=t-5}^{t-1} \exp(\hat{\beta}_a) \mathbb{1}_{\text{age}_{i\tau}=a}}$$

We use this measure to rank workers within each (year) period \times (year-of-birth) cohort cell. Based on this rank, we create 21 recent wages group: workers below the 5th percentile of the recent wages distribution (conditional on age and year), those between the 5th and the 10th percentile of the recent wages distribution, etc., those between the 95th and the 99th percentile of the recent wages distribution and finally those above the 99th percentile of the recent wages distribution. We isolate top earners firstly because underrepresentation of women among top earners may explain a massive share of the aggregate gender gap (Fortin, Bell, and Böhm, 2017).

³An insightful presentation of this method is provided by Afssa and Buffeteau (2006).

4.2 Labor earnings decomposition

Both labor earnings – defined as the sum of all salaried earnings over a given year – and working time – in paid hours of work – are available in our dataset. Let denote labor earnings (resp. working time) of individual i during year t as \tilde{y}_{it} (resp. \tilde{l}_{it}). We consider a normalized version of earnings (working time) net of age effect. To do so, we regress log-labor earnings (log-working time) on a set of age, period and cohort dummies similar as 1. Hourly wages being defined as the ratio of annual earnings over working time implies that: $\beta_a^y = \beta_a^w + \beta_a^l$. We use the estimated coefficients to define normalized current earnings, wages and working time: $y_{it} = \tilde{y}_{it} / \exp(\hat{\beta}_a^y) =$, $w_{it} = \tilde{w}_{it} / \exp(\hat{\beta}_a^w) =$ and $l_{it} = \tilde{l}_{it} / \exp(\hat{\beta}_a^l)$.

We now consider average earnings at time $t + k$ ($k = 0 \dots 10$) for a group of workers defined by their rank in the distribution of recent wages. We include individuals with labor earnings inferior to $1/8 \underline{w}_{t+k}$ in $t + k$ as having labor earnings equal to 0. Let $d_{i,t+k}$ denote a dummy of having earnings superior to this threshold in $t + k$. Then $y_{i,t+k} = d_{i,t+k} l_{i,t+k} w_{i,t+k}$. Hence we can decompose average labor earnings at time $t + k$:

$$\begin{aligned} \log(\mathbb{E}[y_{i,t+k}]) &= \log(\mathbb{P}(d_{i,t+k} = 1)) \\ &+ \log(\mathbb{E}[l_{i,t+k} | d_{i,t+k} = 1]) \\ &+ \log\left(\frac{\mathbb{E}[w_{i,t+k} l_{i,t+k} | d_{i,t+k} = 1]}{\mathbb{E}[l_{i,t+k} | d_{i,t+k} = 1]}\right) \end{aligned} \tag{2}$$

The first term corresponds to average labor market participation at year $t + k$, the second represents average working time for those who participate, and the last one is average hourly wages for participants, weighted by their working time. This decomposition is easily adapted to the gender gap in average earnings at time $t + k$, allowing to disentangle differences in participation, working time and hourly wages among groups of workers that have very similar hourly wages between $t - 5$ and $t - 1$.

Practically speaking, we implement this decomposition within each recent wages \times 5-year age group (23-24, 25-29, 30-34, ..., 50-54) cell and average over all age groups.

4.3 Consequences of career interruptions and wage growth

In order to assess whether career interruptions have different consequences for men and women, and whether men and women who do not experience career interruptions may have different wage growth, we compare them in $t + k$ ($k = 1 \dots 10$) and separate those that experienced a full year of nonemployment in t (non-participants) from those who did not leave employment in t .

Among non-participants, we use the same decomposition as 2. Among participants, we consider another decomposition:

$$\begin{aligned} \log(\mathbb{E}[y_{i,t+k} | d_{it} = 1]) &= \log(\mathbb{P}(d_{i,t+k} = 1 | d_{it} = 1)) \\ &+ \log(\mathbb{E}[l_{i,t+k} | d_{it} = 1, d_{i,t+k} = 1]) \\ &+ \log\left(\frac{\mathbb{E}[w_{i,t} l_{i,t+k} | d_{it} = 1, d_{i,t+k} = 1]}{\mathbb{E}[l_{i,t+k} | d_{it} = 1, d_{i,t+k} = 1]}\right) \\ &+ \log\left(\frac{\mathbb{E}\left[\frac{w_{i,t+k}}{w_{i,t}} w_{i,t} l_{i,t+k} | d_{it} = 1, d_{i,t+k} = 1\right]}{\mathbb{E}[w_{i,t} l_{i,t+k} | d_{it} = 1, d_{i,t+k} = 1]}\right) \end{aligned} \tag{3}$$

The first two terms have the same interpretation as those of 2. The third term corresponds to average hourly wages at time t , weighted by current working time at time $t + k$. The last

one refers to hourly wage growth, weighted by counterfactual earnings if individuals did not experience any wage growth between t and $t + k$ but had the same working time at time $t + k$.

4.4 Consequences of childbirth

Recent research has highlighted how childbirth, and specifically first childbirth, may shift labor market trajectories of women with respect to those of men (Kleven, Landais, and Søgaard, 2018). Missing data from birth records during the 1990s makes it difficult to study what happens at the time of the first childbirth, or to isolate individuals that never had children. Because of this, we simply compare at time $t + k$ ($k = 0...10$) individuals that had children during year t with those who did not have children between t and $t + k$ (but may have had children before t or after $t + k$).

4.5 Distributions of wage growth

In a last part, we dig further into differences of wage growth between men and women. In order to do so, we focus on individuals that participate in the labor market both at time t and at time $t + k$, and investigate for gender differences in wage growth very similar to Weinberger (2011). For those workers, we measure (normalized) wage growth as $\delta^k w_{it} = \log(w_{i,t+k}) - \log(w_{it})$. We assess whether women have slower wage growth than men, or are penalized by childbirth by estimating:

$$\begin{aligned} \delta^k w_{it} &= \zeta \text{ female}_i \\ &+ \eta (1 - \text{female}_i) \text{ childbirth}_{i,t+1} \\ &+ \kappa \text{ female}_i \text{ childbirth}_{i,t+1} \\ &+ \lambda X_{i,t} + \nu_{it} \end{aligned} \quad (4)$$

where female_i is a gender dummy, $\text{childbirth}_{i,t+k}$ stands for experiencing childbirth in $t + 1$, X_{it} is a set of conditioning variables and ν_{it} is an idiosyncratic error term with mean 0 and finite variance.

Our set of controls includes period dummies, a quartic in age, recent experience defined as the sum of paid hours of work between t and $t + k$, recent experience interacted with age in order to allow for diminishing returns, education, industry (1-digit level), and a dummy for career interruptions defined as having experienced a full-year of nonemployment between t and $t + k$.⁴ We estimate this model separately for each recent wage cell.

Lastly, because these gender differences may vary depending on whether we focus on individuals that experience favorable or unfavorable wage changes (Weinberger, 2011), we estimate similar quantile regression models (Koenker and Bassett, 1978) for different quantile orders τ :

$$\begin{aligned} \mathcal{Q}_\tau(\delta^k w_{it} | \text{female}_i, \text{childbirth}_{i,t+1}, X_{it}) &= \zeta_\tau \text{ female}_i \\ &+ \eta_\tau (1 - \text{female}_i) \text{ childbirth}_{i,t+1} \\ &+ \kappa_\tau \text{ female}_i \text{ childbirth}_{i,t+1} \\ &+ \lambda_\tau X_{i,t} \end{aligned} \quad (5)$$

Standard errors and confidence intervals are estimated by subsampling, which is computationally attractive given large sample size (Chernozhukov and Fernández-Val, 2005; Arellano and Bonhomme, 2017). We choose subsampling size as a constant plus the square root of the sample size, where the constant (4000) was taken to ensure reasonable finite sample performance of the estimator.

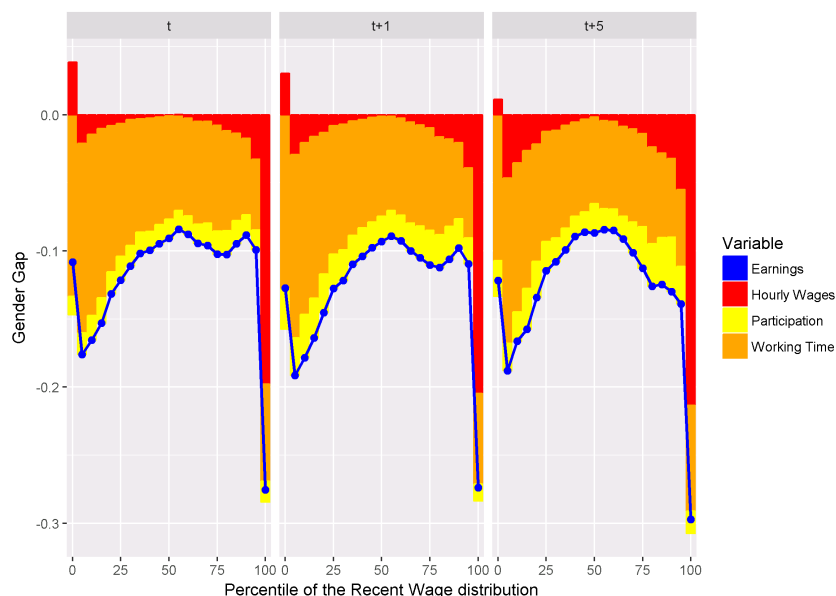
⁴only when $k > 1$.

5 Results

5.1 Decomposition of the gender gap in earnings

We start by comparing annual earnings of men and women conditional on rank in the distribution of recent wages (see Subsection 4.2). Figure 1 displays the decomposition of the log-gender gap in average annual earnings (including 0 for individuals that do not participate in the labor market) between differences in participation rate, differences in working time and differences in hourly wages, depending on their position in the hourly wage distribution between $t - 5$ and $t - 1$.

Figure 1: Average gender gap in t , $t + 1$ and $t + 5$ by position in the distribution of recent wages



With the exception of both ends of the distribution, we find that the gender gap in earnings at time t is a decreasing function of rank in the recent wage distribution: it varies between 18 log-points for workers between the 5th and the 10th percentiles and 10 log-points for those between the 95th and 99th percentiles. This is mostly because differences in working time (conditional on labor market participation) decrease with recent wages: from 14 log-points for the first group to 5 log-points for the second. The gap in participation varies very little along the distribution: it is 1.6 log-points for the former group of workers and 1.4 log-points for the latter.

The gender gap in hourly wages exhibits a U shape along the recent wages distribution: it is 2 log-points for workers between the 5th and the 10th percentiles, less than 0.1 log points for median workers and 1.8 log-points among workers belonging to the P95-P99 recent wage bin. This U-shape intensifies as times goes by: at time $t + 5$, the gender gap in hourly wages is 4.6 log-points for the first group of workers, 0.2 log-points for the second group and 5.5 log-points for the last group. This results in the gender gap in earnings getting more U-shaped as time passes.

At the lower end of the recent wage distribution, the gender gap in earnings is smaller than it is for workers placed slightly upper in the distribution. This is mostly because in this group, women tend to have better hourly wages than their male counterparts: the difference is 3.9 log-points at time t and decreases to 1.1 log-points at time $t + 5$.

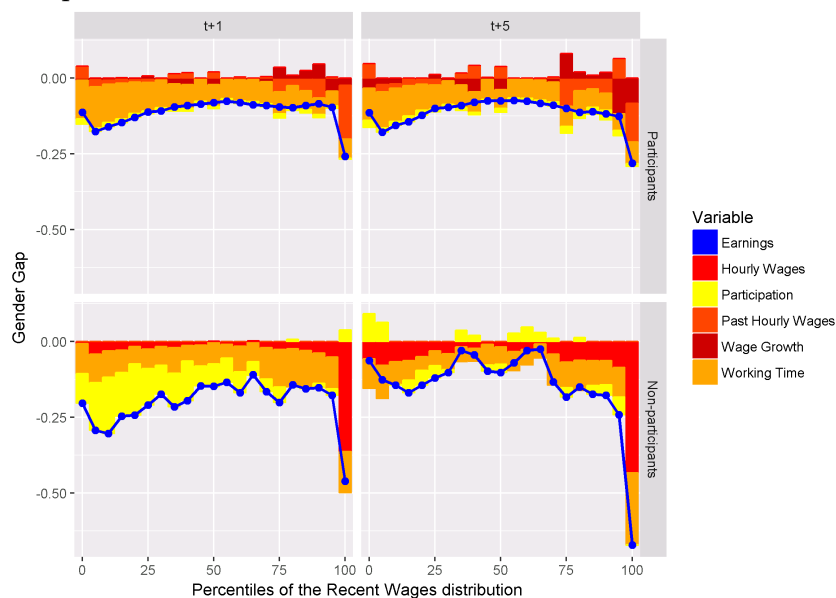
Among top-earners (defined as the top 1% of the recent wage distribution), the gender gap in earnings is larger than it is for every other group of workers: it is 28 log-points at time t . Differences in hourly wages explain the largest part of the gap: women have hourly wages inferior

by 19.8 log-points to those of their male counterparts. The difference increases slightly over time, to 30 log-points at time $t + 5$ for the gender gap in earnings and 21.3 log-points for the gender gap in hourly wages.

5.2 Career interruptions and wage growth for participants

In order to assess whether those differences arise from career interruptions, initial heterogeneity within each of the recent wages cells or differences in wage growth rate, we isolate workers that experienced a full year of nonemployment at time t and observe their earnings, labor market participation, working time and hourly wages at time $t + k$. For those who participated in the labor market at time t , we decompose their hourly wages at time $t + k$ between hourly wages at time t and hourly wages growth between t and $t + k$. Figure 2 displays our results.

Figure 2: Average gender gap in $t + 1$ and $t + 5$ by position in the distribution of recent wages and participation in t



Among participants, and with the exception of top-earners, we find that the gender gap in earnings is primarily driven by differences in working time. Consistent with our previous findings, these differences are larger among low-productivity workers than among those with higher hourly wages. For the larger, lower part of the recent wage distribution, the gap in hourly wages remains small, and largely explained by differences in past hourly wages rather than differences in hourly wages growth. In the upper half, with the exception of top earners, differences in hourly wages are small too, but rather because differences in past hourly wages (detrimental to women) are compensated by differences in hourly wages growth (detrimental to men).


Among top earners, we find a massive gender gap in hourly wages, consistent with that of Figure 1. It amounts to 20.1 log-points at time $t + 1$ and 21.0 log-points at time $t + 5$. This implies that this huge difference is not likely the mere result of career interruptions. Differences in initial wages explain the larger part of this considerable wage gap: 17.7 log-points at time $t + 1$ (2.4 log-points for hourly wages growth) and 12.7 log-points (8.3 log-points for hourly wages growth). Hence the gender gap among top-earners is likely the result of considerable wage heterogeneity and strong gender gradient, rather than high-achieving women experiencing much slower wage growth. Evidence from wage growth is however not clearcut; in Section 5.4 we investigate it in a more systematic way.

When it comes to non-participants, career interruptions might be more persistent for women than they are for men in the larger, lower part of the recent wages distribution. The gender gap in hourly wages among workers who experienced a full year of nonemployment in t is larger at both ends of the distribution than they are at the middle. Finally, the gender gap in hourly wages among top earners who experienced a full year of nonemployment is very large (36.4 log-points at time $t + 1$). This might be either because of preexisting differences between men and women in the wage rate, which would be consistent with the pattern among participants, or because women have more difficult recoveries than men do.

5.3 Consequences of childbirth

We investigate the contribution of motherhood to the gender pay gap by comparing average earnings, participation, working time and hourly wages at time $t + k$ for workers that had a child during year t with those of workers that did not have any child between t and $t + k$. Figure 3 displays our results, separately for men and women, for each position in the recent wages distribution

Figure 3: Average parenthood gap in t , $t + 1$ and $t + 5$ by position in the distribution of recent wages and gender



Among men, we find recent fathers to have earnings, participation, working time and hourly wages that are very similar from those of their counterparts that did not have any child between t and $t + k$, with the exception of both ends of the distribution. Among very low wage earners and top earners, recent fathers have higher earnings: by 4.5 log-points at time t for the first group and 13.4 log-points for the second group. The difference is largely explained by differences in hourly wages, which accounts for 6.0 log-points for the former and 8.9 log-points for the latter. As time goes by, the gap between recent fathers and their counterparts widens, so that the former tend to have higher earnings and hourly wages than the latter; the difference is larger at both ends of the distribution.


When it comes to women, childbirth correlates with massive and very heterogeneous earnings differences. Specifically, we find recent mothers to have much lower earnings than their counterparts who did not have children between t and $t + k$, but the difference is very large among low wage women and much smaller (and may actually reverse itself) among top earning women. At

time t , new mothers have earnings inferior by 75 log-points to those of women that did not have any child during year t when they belong to the lowest ranks of the recent wage distribution; among female top earners, the difference shrinks to 5.8 log-points.

Differences in participation and working time largely explain the pattern: while among high-achieving women, childbirth correlates with small differences in participation and working time, among their low productivity counterparts the correlation is massive. Hourly wage differences exhibit a quite different pattern. Namely, we find that the very year of childbirth, women may have substantially lower hourly wages than their counterparts who did not give birth, the difference being maximal around the 70th percentile of the recent wages distribution (17.8 log-points) and virtually nonexistent at both ends of the recent wages distribution. As time passes, this motherhood wage gap diminishes and gets approximately uniform along the recent wage distribution: it fluctuates between 4 and 5 log-points at time $t + 1$ and 3 and 4 log-points at time $t + 5$. Both ends of the distribution are again an exception: the gap reverses itself, so that women who gave birth have higher hourly wages than those who did not; the difference amounts to 5.9 log-points among female top earners at time $t + 5$. Note that these results somehow depart from those of Kleven, Landais, and Sogaard (2018) who find the motherhood gap in earnings to be approximately equal to 20% in the long-run, being explained in roughly equal proportions participation, working time and hourly wages.

Figure 4 provides another view of the same data, this time comparing earnings, participation, working time and hourly wages for men and women, separately for those who had a child during year t and those who did not have any child between t and $t + k$.

Figure 4: Average gender gap in t , $t + 1$ and $t + 5$ by position in the distribution of recent wages and childbirth in t



Consistent with our previous findings, we find the gender gap to be much larger among recent parents than it is among individuals that did not have any child between t and $t + k$. This is mostly because recent mothers have lower working time than women who did not recently give birth. However, we also find that the gap in hourly wage is larger among recent parents than it is among their counterparts that do not experience recent childbirth.

However, we find that even among top earners who did not have any child between t and $t + k$, the gender gap in hourly wages is very large. It amounts to 20 log-points in t (21 log-points in $t + 1$ and $t + 5$), against 28 log-points among recent parents (29 log-points in $t + 1$ and $t + 5$).

While the latter being substantially larger than the former may imply that motherhood penalties do indeed contribute to the massive gender gap among top earners (28 log-points overall, see Figure 1), the magnitude of the gap among individuals that did not experience recent childbirth makes it unlikely that it is the main channel.

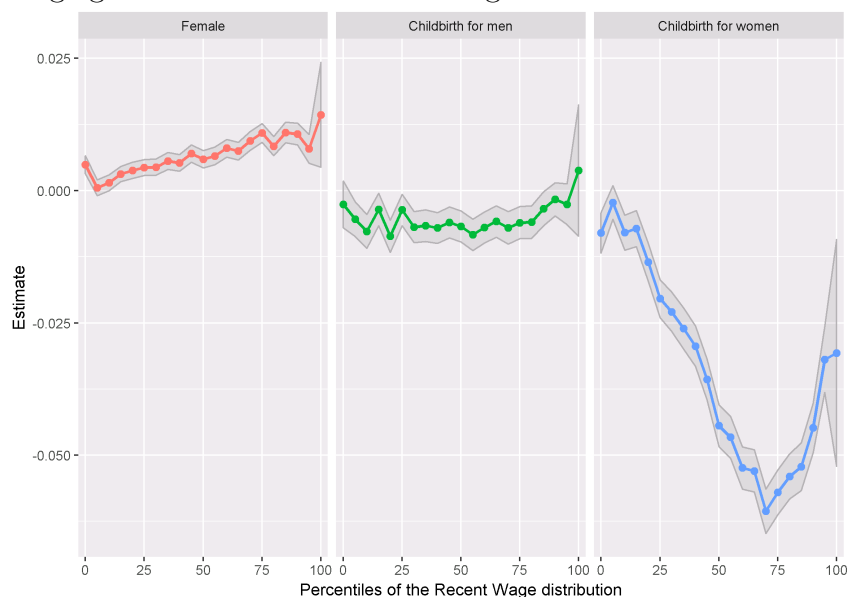
While missing data means we do not isolate individuals who never had child from those who did, making the evidence not as clearcut as we wish, in Appendix A we display similar results disaggregated by age (see Figure A.2 and A.3). We find massive a massive gender gap in hourly wages among top earners aged 23-24 who did not experience recent childbirth. Since those are plausibly without children, we consider it evidence that large gender wage gap among top-earners is not merely the result of childbirth.

5.4 Evidence from the distribution of wage growth

5.4.1 Gender and motherhood gap

We finally turn to evidence from wage growth (conditional on continued labor market participation), which allows us to abstract from potential heterogeneity in initial hourly wages. In order to do so, we impliment separate regression of individual annual hourly wages changes $\delta^1 w_{it} = \log(w_{i,t+1}) - \log(w_{it})$ on gender and childbirth (plus a set of additional controls: see Subsection 4.5). Figure 5 displays our OLS estimates.

Figure 5: Wage growth between t and $t + 1$: gender and childbirth OLS estimates




We find that in average, women with continued participation that do not experience childbirth do not have slower wage growth than their male counterparts. Actually, with the exception of the lower end of the distribution, we find that they may have significantly better wage growth; this female premium would amount to 1.5 log-points for female top earners. This may however reflect mothers catching up with other women after childbirth.

Recent fathers may experience slightly slower wage growth than their counterparts who did not have new child in $t + 1$, except at both ends of the recent wages distribution where we do not find any significant differences. However, the difference remains quite small. It is maximal for men between the 20th and the 25th percentiles of the recent wages distribution for which amounts to 1.1 log-points.

Lastly, women who gave birth in $t+1$ experience much lower wage growth in average between t and $t+1$ than their counterparts who did not. The difference is highly heterogeneous along the recent wages distribution. This motherhood penalty exhibits a clear U-shape pattern, similar as our previous findings (see Figure 3). It is minimal at the bottom of the distribution (1.2 log-points) and maximal for women between the 70th and the 75th percentiles of the recent wage distribution (6.4 log-points). It is significantly lower among top earning women (5.1 log-points) even though the confidence intervals are large.

We assess the possibility that these penalties and premium may vary depending on how favorable hourly changes are (conditional on observables) by implementing similar quantile regression for different quantile orders. Figure 6 displays our estimates.

Figure 6: Wage growth between t and $t+1$: gender and childbirth quantile estimates



At the top of the recent wages distribution, our results suggest that women who did not experience childbirth do not suffer more from unfavorable wage shocks than their male counterparts. When they experience favorable shocks, they are indeed likely to experience faster wage growth. Those two findings are inconsistent with the existence of both a glass ceiling and a paper floor (Güvenen, Kaplan, and Song, 2014). Reversely, among low wage earners, we find women to experience less negative wage changes than their male counterparts when exposed to unfavorable shocks, whereas they seem to make less of favorable shocks. This pattern would be consistent with the existence of a sticky floor.

Consistent with our previous findings, men seem very slightly affected by recent childbirth. Reversely, among women childbirth correlates with very large differences in wage growth. Specifically, our estimates suggest that childbirth is associated with large dispersion in wage growth.

This is because the most unfavorable changes experienced by women that gave birth in $t + 1$ are way more negative than those of their counterparts that did not, whereas the most favorable changes that they experience are actually more positive. The size of the former effect is substantially larger than that of the latter, which is consistent with the average effect of childbirth being negative. Those sizes also vary along the recent wages distribution: they are maximal around the 70th-75th percentiles of the distribution, and minimal at both ends.

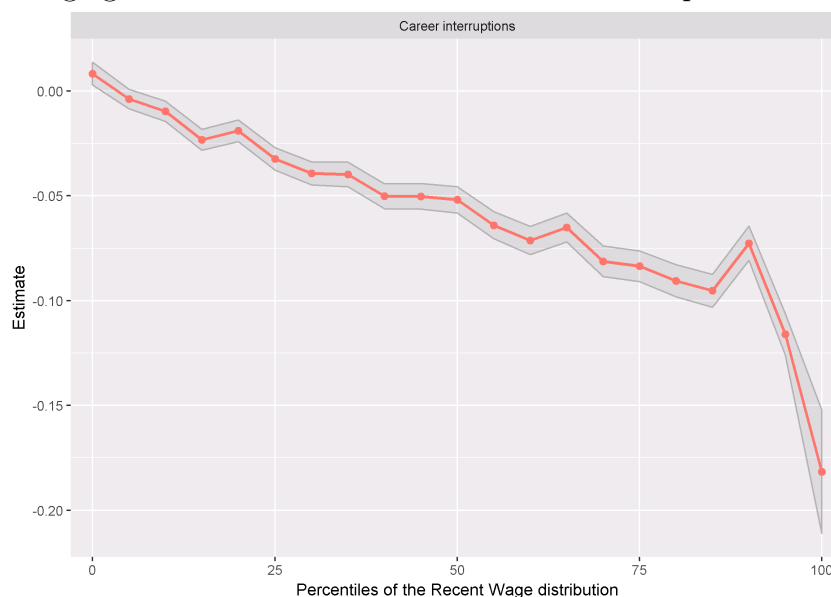
An important point is that while we find top earning women to be affected by childbirth, the estimated sizes of the effect do not depart from the rest of the distribution. Moreover, when for those who do not experience childbirth, we do not find they wage growth to be significantly slower than those of their male counterparts. This contrasts with our evidence that the gender gap in hourly wages among top wage earners is much larger than it is for the rest of the recent wages distribution. This discrepancy suggests that large gender differences in hourly wages among top earners arise at labor market entry rather than along the course of a career.

5.4.2 Cost of career interruptions

However, this does not necessarily mean that those two stylised facts are totally unrelated. Indeed, if top jobs have lower children related amenities, or exhibit higher losses for time spent out of the labor market, then some high-achieving women may likely choose other occupations as they enter the labor market (Adda, Dustmann, and Stevens, 2017). To investigate this possibility, we estimate similar regressions, this time with 5-year wage growth as outcome. Figure 7 plots our estimates of the coefficient of a dummy for having spent at least on year out of the labor market between t and $t + 5$.

This coefficient does not correspond to the true cost of career interruptions. Firstly, our dummy aggregates true career interruptions, i.e. full years spent out of the labor market, with self-employment, employment in the public sector or even international mobilities. Secondly, the choice of interrupting one's career is likely to be endogeneous with respect to anticipated career prospects. Yet we assume heterogeneity in the coefficient along the recent wages distribution can reflect heterogeneity in terms of the cost of career interruptions. This assumption is somehow correct to the extent that our measurement error and endogeneity issues do not vary too much all along the recent wage distribution.

Figure 7: Wage growth between t and $t + 5$: career interruptions OLS estimates



We find the wage penalties for workers who experienced career interruptions between t and $t + 5$ to increase over the recent wage distribution. Namely, the coefficient does not differ significantly from 0 for workers placed lowest in the distribution. It decreases almost linearly – i.e. penalties increase – up until the 90th percentile where it amounts to approximately 9 log-points. There is a break between the 90th and the 95th percentiles, and the slope gets steepest. In the end, top earners experience penalties that are significantly larger than the rest of the workers, including those placed between the 95th and the 99th percentiles of the recent wage distribution. For top earners, the career interruption penalty is approximately 17 log-points. In Appendix A, we display our quantile regression estimates that show substantial heterogeneity, possibly related to the motives of this time spent out of the labor market or heterogeneity in the duration of these career interruptions. However, they consistently show that top earners can experience much larger wage losses from career interruptions than the rest of the workers.

It is therefore possible that women may choose not to enter top jobs as they begin their careers, as they anticipate the cost of time spent out of the labor market to be much larger than it is for other positions. Assessing the contribution of this effect to the underrepresentation of women at the highest end of the wage distribution would likely require a full structural model of career and fertility choices (Adda, Dustmann, and Stevens, 2017) that is beyond the scope of this paper.

6 Conclusion

This investigation of the gender pay gap build on a descriptive, non-parametric framework. Within groups of workers defined by their past recent wages, we find the gender gap in earnings to be explained to a large extent by differences in working time, while the contribution of hourly wages is much smaller. The top of the wage distribution departs from this pattern: among recent top wage earners, women have much lower hourly wages than their male counterparts. This suggests that top inequality and underrepresentation of women at the very top of the wage distribution may indeed contribute to a large extent to the gender pay gap, which is consistent with previous evidence (Fortin, Bell, and Böhm, 2017).

We dig further in the data to confront this pattern with other plausible channels of the gender pay gap. While career interruptions and childbirth, and more generally differences in wage growth as opposed to wage levels may well contribute to the gender wage gap, our results suggest that the effects are not large enough to account for the massive gap among top earners. This offers ground to the claim that high-achieving women not making it to the top of the wage distribution at the same rate as their male counterparts is primarily the result of pre-labor market differences or patterns that arise at labor market entry. Our result suggest that the career costs of time spent out of the labor market are much higher for top earners than they are for the rest of the workers. If high-achieving women anticipate this fact before they enter the labor market, this can cause those of them with the highest taste for children to chose positions that do not belong to the top of the wage distribution.

This does not rule out other channels that contribute to the gender gap getting larger over the years of a career. Career interruptions and childbirth do indeed correlate with increase in the gender pay gap for the vast majority of workers. More precisely, our results suggest, in accordance with recent research (Kleven, Landais, and Søgaaard, 2018), that childbirth explains to a large extent why women tend to have slower wage growth than their male counterparts.

These findings offer practical implications for gender equality-oriented public policies. They suggest that targetting gender differences among high-achieving workers in event that happen over the course of a career, like promotions, is likely less efficient than targetting pre-labor or differences at labor market entry. This may explain why policies aimed at increasing promotions for high-achieving women, like gender quotas, have been found to generate disappointing results (Bertrand et al., 2014; Fortin, Bell, and Böhm, 2017). Indeed, if gender differences among top

earners mostly result from choices made at labor market entry based on anticipated careers, the effects of such policies will only be seen after a new generation of workers replaces those that were primarily targeted. If those anticipations are mediated by educational choices, such as choices of college major (Zafar, 2013), recent evidence suggests ways it can be acted upon (Breda et al., 2018).

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.
- Arellano, M., and S. Bonhomme. 2017. “Quantile Selection Models With an Application to Understanding Changes in Wage Inequality.” *Econometrica* 85:1–28.
- 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.
- Bertrand, M., S.E. Black, S. Jensen, and A. Lleras-Muney. 2014. “Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labor Market Outcomes in Norway.” Working Paper No. 20256, National Bureau of Economic Research.
- 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.
- Breda, T., J. Grenet, M. Monnet, and C.V. Effenterre. 2018. “Can female role models reduce the gender gap in science? Evidence from classroom interventions in French high schools.” Pse working papers, HAL, Feb.
- Bütikofer, A., S. Jensen, and K.G. Salvanes. 2018. “The Role of Parenthood on the Gender Gap among Top Earners.” Discussion Paper Series in Economics No. 9/2018, Norwegian School of Economics, Department of Economics, Apr.
- 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.
- Chernozhukov, V., and I. Fernández-Val. 2005. “Subsampling Inference on Quantile Regression Processes.” *Sankhyā: The Indian Journal of Statistics (2003-2007)* 67:253–276.
- Christofides, L.N., A. Polycarpou, and K. Vrachimis. 2013. “Gender wage gaps, ‘sticky floors’ and ‘glass ceilings’ in Europe.” *Labour Economics* 21:86 – 102.
- 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.
- 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.
- 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.
- Gardeazabal, J., and A. Ugidos. 2005. "Gender Wage Discrimination at Quantiles." *Journal of Population Economics* 18:165–179.
- 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.
- Goldin, C., S. Pekkala Kerr, C. Olivetti, and E. Barth. 2017. "The Expanding Gender Earnings Gap: Evidence from the LEHD-2000 Census." *American Economic Review* 107:110–14.
- 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.
- Guvenen, F., S. Ozkan, and J. Song. 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy* 122:621–660.
- 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. Søgaaard. 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. 2005. "The evolution of the gender wage gap." *Labour Economics* 12:73 – 97.
- . 2003. "Gender Differences in Entry Wages and Early Career Wages." *Annales d'Économie et de Statistique*, pp. 245–265.
- Manning, A., and H. Robinson. 2004. "Something in the way she moves: a fresh look at an old gap." *Oxford Economic Papers* 56:169–188.

- Manning, A., and J. Swaffield. 2008. "The Gender Gap in Early-Career Wage Growth." *The Economic Journal* 118:983–1024.
- 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.
- Morgan, L.A. 1998. "Glass-Ceiling Effect or Cohort Effect? A Longitudinal Study of the Gender Earnings Gap for Engineers, 1982 to 1989." *American Sociological Review* 63:479–493.
- 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.
- 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.
- Weinberger, C.J. 2011. "In Search Of The Glass Ceiling: Gender and Earnings Growth among U.S. College Graduates in the 1990s." *Industrial and Labor Relations Review* 64:949–980.
- Weinberger, C.J., and P.J. Kuhn. 2010. "Changing Levels Or Changing Slopes? The Narrowing of the Gender Earnings Gap 1959 - 1999." *Industrial and Labor Relations Review* 63:384–406.
- Wilner, L. 2016. "Worker-firm matching and the parenthood pay gap: Evidence from linked employer-employee data." *Journal of Population Economics* 29:991–1023.
- Zafar, B. 2013. "College Major Choice and the Gender Gap." *Journal of Human Resources* 48:545–595.

A Additional results

Figure A.1: Average gender gap in t , $t + 1$ and $t + 5$ by position in the distribution of recent wages and age

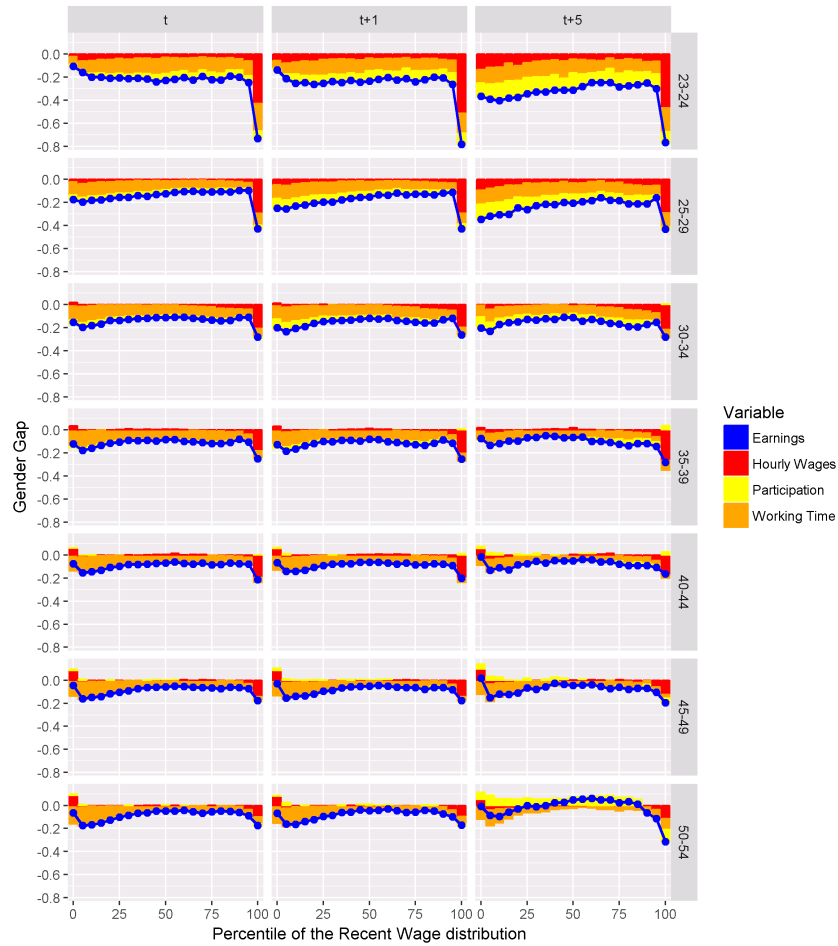


Figure A.2: Average gender gap in t , $t + 1$ and $t + 5$ by position in the distribution of recent wages and age for recent parents

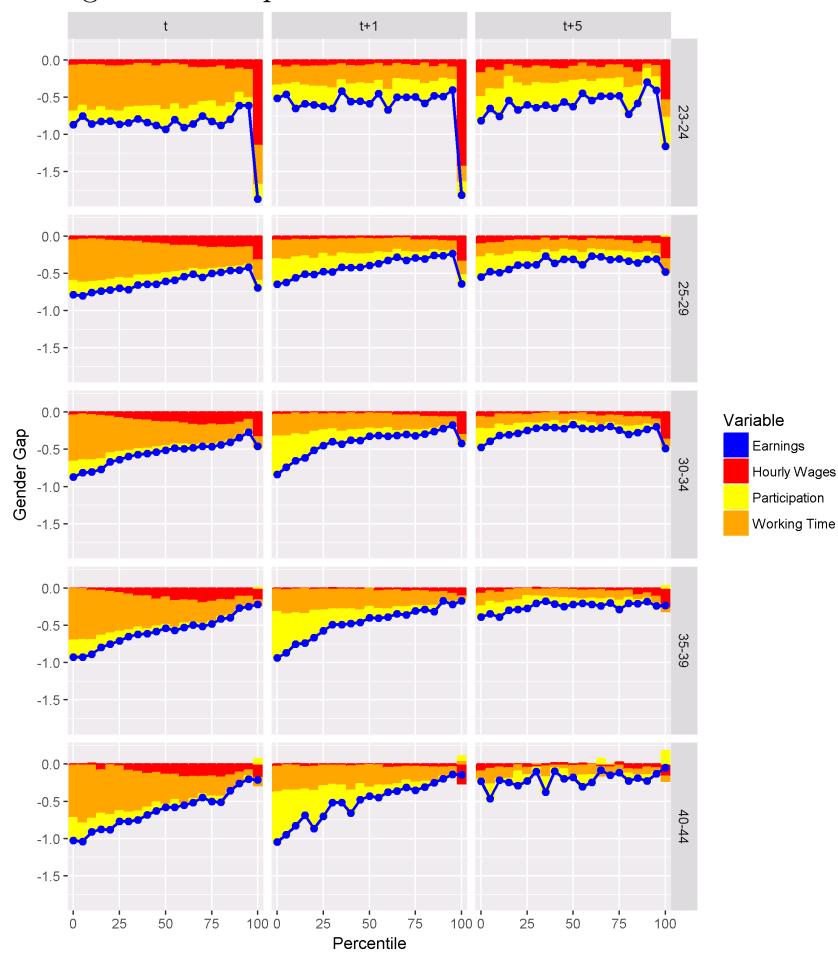


Figure A.3: Average gender gap in $t + 1$ and $t + 5$ by position in the distribution of recent wages and age for non-recent parents

