



HAL
open science

Gender, family and labor, supply

Pierre Pora

► **To cite this version:**

Pierre Pora. Gender, family and labor, supply. Economics and Finance. Université de Nanterre - Paris X, 2022. English. NNT : 2022PA100007 . tel-03817180

HAL Id: tel-03817180

<https://theses.hal.science/tel-03817180>

Submitted on 17 Oct 2022

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

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

Membre de l'université Paris Lumières

Pierre Pora

Gender, Family and Labor Supply

Thèse présentée et soutenue publiquement le 6 janvier 2022
en vue de l'obtention du doctorat de Sciences économiques
de l'Université Paris Nanterre
sous la direction de Mme Dominique Meurs (Université Paris Nanterre)

Jury :

Rapporteur-e :	M. Bertrand Garbinti	Professor, Ensaе-Crest
Rapporteur-e :	Mme Anne Solaz	Directrice de recherche, Ined
Membre du jury :	M. Éric Maurin	Directeur d'études, EHESS-PSE
Membre du jury :	M. Roland Rathelot	Professor, Ensaе-Crest

Remerciements

Je tiens tout d'abord à remercier Dominique Meurs d'avoir accepté il y a trois ans de diriger cette thèse, et de m'avoir permis de la mener dans les meilleures conditions possibles. Ses conseils et suggestions ont nettement amélioré ma production sur le plan scientifique, en me sortant à l'occasion d'impasses dans lesquelles j'avais bien entendu foncé tête baissée. Son soutien dans les moments de doute m'a aussi été très précieux.

J'adresse ensuite mes remerciements à Bertrand Garbinti, Éric Maurin, Roland Rathelot et Anne Solaz d'avoir accepté de faire partie de mon jury. Mon travail a largement bénéficié de leur relecture et de leurs remarques. Je suis particulièrement reconnaissant à Bertrand Garbinti et Anne Solaz qui ont accepté d'être rapporteurs.

Cette thèse doit évidemment beaucoup à mes co-auteurs, Raphaël Lardeux et Lionel Wilner. J'ai énormément appris de nos collaborations, et sur tous les plans. Je remercie également les co-auteurs avec lesquels j'ai pu travailler au cours des dernières années sur des projets qui, pour n'avoir pas trouvé leur place dans cette thèse, ont néanmoins largement nourri ma réflexion sur certaines des questions qu'elle aborde : Odran Bonnet, Simon Georges-Kot et Mélina Hillion.

Je remercie tous ceux qui ont contribué à rendre ma présence à l'Université Paris Nanterre, et en particulier à EconomiX, aussi agréable que possible, bien qu'elle soit restée largement virtuelle. En particulier, je remercie Fanny Claise, Jérôme Deyris, María José Montoya pour leur aide dans l'organisation de mes interventions en séminaire, ainsi que Magali Dumontet, Paul Malliet, Benjamin Monnery et Julie Tréguier pour leurs discussions très enrichissantes. Je suis bien entendu extrêmement reconnaissant à Frédéric Hammerer sans lequel je n'aurais pas été capable de venir à bout des formalités administratives nécessaires à l'aboutissement de cette thèse.

Cette thèse est le produit d'une trajectoire de plusieurs années au sein du service statistique public. Je remercie en premier lieu Élise Coudin, Emmanuel Berger, Laurence Rioux et Lionel Wilner, pour m'avoir fait entrer, à ma sortie de l'Ensaë, dans l'univers des données administratives sur le marché du travail qui fournissent l'essentiel de sa matière empirique à cette thèse. Je dois encore remercier Élise Coudin et Lionel Wilner une seconde fois pour m'avoir poussé à m'inscrire en thèse tant que mon poste à l'Insee m'en laissait le temps ; j'aurais très vraisemblablement

encore repoussé l'échéance sans leurs encouragements.

Je remercie vivement Dominique Goux et Sébastien Roux pour m'avoir accueilli au Département des Études Économiques de l'Insee pendant trois années au cours desquelles j'ai énormément appris. Les deux premiers chapitres de cette thèse doivent également beaucoup à leurs relectures attentives. Mes remerciements s'adressent aussi à tous les chargés d'études que j'ai côtoyés : tout d'abord Rémi Monin, Milena Suarez-Castillo et Cécile Welter-Médée qui ont cohabité avec moi dans le bureau 4B296, mais aussi Mathias André, Arthur Bauer, Josselyn Boussard, Simon Bunel, Lino Galiana, Olivier Meslin, Simon Quantin, Julien Silhol et plus largement tous ceux qui m'ont entouré pendant ces trois années. Je remercie aussi Isabelle de Beaumont et Brigitte Rigot qui ont largement contribué à aplanir toutes les difficultés matérielles que j'ai pu rencontrer.

Je remercie chaleureusement Muriel Barlet, Hélène Chaput, Quentin Laffeter et Noémie Vergier pour m'avoir non seulement offert les conditions de travail qui m'ont permis d'achever cette thèse en poste à la sous-direction de l'Observation de la Santé et de l'Assurance Maladie de la Drees, mais encore la possibilité d'approfondir suffisamment les sujets qui sont aujourd'hui les miens pour alimenter les deux derniers chapitres. Mes remerciements vont aussi à tous les chargés d'études du Bureau des Professions de Santé pour leur soutien tout au long de l'année qui s'est écoulée.

J'ai pu tout au long de l'écriture de cette thèse, et plus largement depuis mon arrivée à l'Insee en 2014, bénéficier d'une affiliation au Crest. Je remercie Francis Kramarz et Ivaylo Petev de m'avoir permis de renouveler cette affiliation tout au long de ces années. Je remercie également Sander Wagner pour nos échanges récents qui me poussent à souhaiter poursuivre ces travaux sur l'interaction entre vie familiale et vie professionnelle, ainsi que Philippe Choné et Xavier D'Haultfœuille pour les encouragements qu'ils m'ont prodigués lors de nos rencontres.

Mes premiers pas dans la recherche en sciences sociales quantitatives, incertains et maladroits, ont été effectués grâce à Olivier Godechot. Je le remercie pour son soutien répété depuis, et je lui demande de m'excuser d'être finalement devenu une sorte d'économiste.

La liste de ceux qui, par une question, une remarque ou une suggestion, ont contribué à améliorer ma production est trop longue pour pouvoir tous les énumérer. Je leur demande de me pardonner de ne pouvoir mentionner leurs noms ici, et je les

remercie encore d'avoir partagé avec moi les réflexions que pouvaient susciter chez eux les questions auxquelles je me suis consacré.

Il me faut enfin remercier tous ceux, trop nombreux pour être cités, qui, amis ou membres de la famille, m'ont entouré, en particulier au long d'une année 2021 éprouvante. Je pense en particulier à mes parents et à mon frère, toujours présents pour moi. Je ne peux finir ces remerciements sans une pensée pour Anne, qui se moquerait probablement de moi pour soutenir à un âge aussi avancé, mais dont l'absence laisse un immense vide dans nos vies.

Résumé et mots-clés

Cette thèse étudie de façon empirique l'effet de la vie familiale sur l'offre de travail dans le contexte français. Elle comprend quatre chapitres indépendants, chacun consacré à un aspect précis de l'impact de la vie familiale sur les décisions d'allocation de temps au travail à l'extérieur du foyer. Le premier contribue à la littérature récente traitant de l'effet de la fécondité sur les revenus et l'offre de travail des mères en examinant l'hétérogénéité de ces effets le long de la distribution de leurs salaires potentiels. Le second cherche à déterminer si l'augmentation de la disponibilité de services de garde d'enfants abordables fournis par des crèches subventionnées par les pouvoirs publics, et donc la diminution du coût implicite du travail à l'extérieur du domicile, permet aux parents, et en particulier aux mères de jeunes enfants, d'augmenter leur offre de travail et leurs revenus. Le troisième se concentre sur les conséquences de la vie familiale sur les systèmes de santé, en quantifiant la contribution de la parentalité à la diminution de l'offre de travail des infirmières au long de leurs carrières. Enfin, le dernier chapitre explore comment un changement brutal et permanent, ici le licenciement d'un des adultes de la famille, affecte la structure de l'avantage comparatif intra-familial et se répercute sur les relations familiales et l'offre de travail des autres membres.

Mots-clés : Inégalités de genre, offre de travail au sein de la famille, garde d'enfants, travail infirmier, licenciement économique

Abstract and keywords

This dissertation investigates the effect of family life on labor supply in the French context through an empirical lens. It consists of four independent chapters, each devoted to a specific aspect of the impact of family life on decisions to allocate time to work outside the home. The first contributes to the recent literature on the effect of fertility on mothers' earnings and labor supply by assessing the heterogeneity of these effects along the distribution of their potential wages. The second examines whether increasing the availability of affordable childcare provided by publicly subsidized daycare centers, and thus lowering the implicit cost of working outside the home, allows parents, and in particular mothers of young children, to increase their labor supply and earnings. The third focuses on the impact of family life on health care systems, and quantifies the contribution of parenthood to the decline in nurses' labor supply over their careers. Finally, the last chapter explores how a sudden and persistent disruption, in this case the layoff of one of the adults in the family, affects the structure of within-family comparative advantage and spills over into the family relationships and labor supply of the other members.

Keywords: Gender inequality, family labor supply, childcare, nursing, job displacement.

Contents

Remerciements	i
Résumé et mots-clés	iv
Abstract and keywords	v
Introduction	1
Vie familiale et écarts entre les sexes sur le marché du travail	2
L’offre de travail des mères : quelles théories en concurrence ?	5
Structure et objet de la thèse	8
Données et méthodes	9
Résumé des chapitres et contributions	10
Bibliographie	22
1 Dissecting child penalties	25
1.1 Introduction	26
1.2 Data and institutional background	31
1.2.1 Data	31
1.2.2 Summary statistics	33
1.2.3 Institutional background	34
1.3 Empirical analysis	36
1.3.1 Normalization	36
1.3.2 Ranks in the hourly wage distribution	37
1.3.3 Difference-in-difference strategy	38
1.4 Results	41
1.4.1 Heterogeneous consequences of childbirth	41
1.4.2 Long-run child penalties	43

1.5	Threats to identification	43
1.5.1	Endogeneity of fertility decisions	44
1.5.2	Mean reversion	45
1.5.3	Selection into treatment	46
1.5.4	Endogeneity of pre-childbirth wages	47
1.6	Conclusion	49
	References	51
	Figures	55
	Tables	66
1.A	Earnings and working time measures	68
1.A.1	Earnings	68
1.A.2	Days	68
1.A.3	Hours	69
1.B	Childbirth imputation	71
1.C	Age-Period-Cohort models	74
1.D	Accounting decomposition	76
1.E	Right-censoring and measurement error	78
1.F	Alternate ranking	90
1.G	Subsequent fertility decisions	92
1.H	Childcare preferences	94
2	Collective childcare and parental earnings	97
2.1	Introduction	98
2.2	Institutional setting	102
2.2.1	Early childcare coverage	102
2.2.2	EAJE-PSU facilities	103
2.2.3	National expansion plans	104
2.2.4	Parental leave policies	104
2.3	Data	105
2.3.1	Family insurance data	105
2.3.2	Labor market data	106
2.3.3	Fertility data	106
2.3.4	Data preparation	107

2.4	Empirical analysis	108
2.4.1	Granular childcare expansions	108
2.4.2	Event-study analysis	110
2.4.3	Instrumental variable approach	112
2.4.4	Identifying assumptions	113
2.5	Parental earnings and labor supply effects	114
2.5.1	Graphical analysis	114
2.5.2	Instrumental variable estimation	115
2.6	Substitution across childcare solutions	116
2.6.1	Paid parental leave	116
2.6.2	Individualized childcare	117
2.7	Conclusion	118
	References	120
	Figures	124
	Tables	132
2.A	Institutional background	136
2.A.1	EAJE-PSU institutions	136
2.A.2	National plans	138
2.A.3	Parental leave policies	139
2.A.4	Childcare preferences and choices	140
2.B	Data	145
2.B.1	Family insurance data	145
2.B.2	Labor market data	146
2.B.3	Fertility data	151
2.B.4	Sample definition	152
2.C	Treatment groups composition	153
2.D	Policy evaluation	156
2.D.1	Aggregate labor supply effect	156
2.D.2	Impact on public finances	158
2.E	Decomposition of the national annual growth in childcare coverage	161
2.F	Identification	163
2.F.1	Sample selection	163
2.F.2	Childcare shock exogeneity	166

2.F.3	Parental earnings	170
2.F.4	Substitution effects	185
3	Do children explain nurses shortages?	197
3.1	Introduction	198
3.2	Institutional context	200
3.3	Data	201
3.3.1	Labor market data	201
3.3.2	Fertility data	203
3.3.3	Sample construction	204
3.3.4	Summary statistics	205
3.4	Empirical analysis	207
3.4.1	Model and identification	207
3.4.2	Aggregation and estimation	210
3.4.3	Simulation exercise	211
3.5	Results	211
3.5.1	Children-related labor supply decisions	211
3.5.2	Contribution to lifecycle profiles of labor supply	213
3.5.3	Robustness checks	214
3.6	Conclusion	216
	References	218
	Figures	220
	Tables	227
3.A	DADS panel: labor supply measures	229
3.A.1	Hours worked: concept	229
3.A.2	Full-time units conversion	229
3.B	How well does the sample cover the nurse occupation?	232
3.C	Additional results	234
3.C.1	Lifecycle profiles	234
3.C.2	Children-related labor supply decisions	235
3.C.3	Overtime, working conditions and hourly wages	241
3.C.4	Contribution to lifecycle profiles of labor supply	243
3.D	Robustness checks	244

3.D.1	Left-censoring issue	244
3.D.2	Right-censoring issue	247
3.D.3	Hours worked measurement	250
3.D.4	Stable control group	253
4	Job displacement, families and redistribution	257
4.1	Introduction	258
4.2	Data	261
4.2.1	Payroll tax registers	262
4.2.2	Permanent Demographic Sample	262
4.2.3	Job displacement identification and sample construction	266
4.2.4	Summary statistics	268
4.3	Empirical analysis	270
4.3.1	Identifying the consequences of job displacement	270
4.3.2	Quantifying the discouraged worker bias	273
4.4	Results	274
4.4.1	Couple formation and dissolution	274
4.4.2	Fertility decisions	276
4.4.3	Labor supply, taxes and transfers	276
4.4.4	Selection into job displacement	280
4.5	Conclusion	281
	References	283
	Figures	285
	Tables	296
4.A	Institutional Background	301
4.B	Discouraged worker effect	304
4.C	Earnings and employment effects in the DADS panel	310
4.D	Individual-level income loss for displaced workers with spouses	312
4.E	Sample selection	314
4.E.1	Impact of job displacement on sample inclusion	314
4.E.2	Replication on the balanced panel	315
4.E.3	Replication on the balanced panel of spouses	321
4.F	Selection into job displacement	325

4.F.1	Restriction to plant-closure events	325
4.F.2	Comparison with non-displaced workers of downsized plants .	333
	Conclusion	339

List of Figures

1	Part de femmes parmi les nouveaux diplômés de l'enseignement supérieur, 2019	3
2	Estimations de l'effet de la parentalité sur la position des parents sur le marché du travail, France 2002-2015	6
3	Estimations de l'effet de la parentalité sur les écarts de rémunération et d'offre de travail entre femmes et hommes, par position initiale dans la distribution de salaire	12
4	Répartition spatiale de l'évolution entre 2007-2015 de l'offre relative d'accueil collectif de jeunes enfants	16
5	Estimations de l'effet de la fécondité sur l'offre de travail des infirmières hospitalières : offre de travail dans le secteur salarié	18
6	Estimations de l'effet de la perte d'emploi d'un salarié sur les revenus et l'offre de travail de son conjoint	20
1.1	Consequences of sample selection with respect to childbirths	55
1.2	Consequences of childbirth for women's labor outcomes	56
1.3	Consequences of childbirth for men's labor outcomes	57
1.4	Impact of first childbirth on the gender gap in earnings and labor outcomes	58
1.5	Probability of having children (sensitivity to the business cycle)	59
1.6	Probability of losing one's job (sensitivity to firm-level shocks)	60
1.7	Probability of having children (sensitivity to firm-level shocks)	61
1.8	Underlying earnings changes	62
1.9	Probability of having children (by subperiod)	63
1.10	Childcare preferences and hourly wages	64
1.11	Counterfactual child penalties based on the reweighted data	65
1.B.1	Imputation of childbirths for individuals born October 2 and 3	73

1.E.1	Consequences of childbirth for women’s labor outcomes: a comparison with a control group determined at the imputed age of childbirth . . .	80
1.E.2	Consequences of childbirth for men’s labor outcomes: a comparison with a control group determined at the imputed age of childbirth . . .	81
1.E.3	Consequences of childbirth for women’s labor outcomes: restriction to older cohorts that have made complete fertility decisions	82
1.E.4	Consequences of childbirth for men’s labor outcomes: restriction to older cohorts that have made complete fertility decisions	83
1.E.5	Consequences of childbirth for women’s labor outcomes: identification based on the timing of the k th childbirth	84
1.E.6	Consequences of childbirth for men’s labor outcomes: identification based on the timing of the k th childbirth	85
1.E.7	Consequences of childbirth for women’s labor outcomes: restriction to childbirths in the second quarter	86
1.E.8	Consequences of childbirth for men’s labor outcomes: restriction to childbirths in the second quarter	87
1.E.9	Consequences of childbirth for women’s labor outcomes: restriction to childbirths in 2000-2010	88
1.E.10	Consequences of childbirth for men’s labor outcomes: restriction to childbirths in 2000-2010	89
1.F.1	Consequences of childbirth for women’s labor outcomes: alternate ranking	90
1.F.2	Consequences of childbirth for men’s labor outcomes: alternate ranking	91
1.G.1	Probability of having a 2nd child along the wage distribution	93
2.1	Childcare prices along the income distribution	124
2.2	Relative supply EAJE-PSU affordable collective childcare at the national level from 2007 to 2015	125
2.3	Spatial distribution of the 2007-2015 growth in relative supply of EAJE-PSU affordable collective childcare	126
2.4	Distribution of maximum annual within-municipality growth in affordable collective childcare coverage	127
2.5	Relative supply of EAJE-PSU affordable childcare, by treatment group and timing of the childcare shock	128

2.6	Event-study estimates of the impact of the childcare shock on mothers' labor earnings, by treatment group	129
2.7	Event-study estimates of the impact of the childcare shock on paid parental leave take-up, by treatment group	130
2.8	Event-study estimates of the impact of the childcare shock on the supply of individualized childcare, by treatment group	131
2.A.1	Ideal childcare solution reporting by parents of children under age 3	141
2.A.2	Actual childcare choices of parents of children under age 3	142
2.A.3	Difficulties in childcare access as reported by parents of children under age 3	143
2.A.4	Impact of childcare availability on labor supply decisions, as reported by parents of children under age 3	144
2.B.1	Imputation of the municipality of residence for jobless observations	150
2.D.1	Decomposition of national-level annual growth in relative supply of EAJE-PSU affordable collective childcare from 2008 to 2014	157
2.F.1	Event-study estimates based on Sun and Abraham (2020) of the impact of the childcare shock on mothers' labor earnings, by treatment group	181
2.F.2	Event-study estimates of the impact of the childcare shock on paid parental leave take-up, by treatment group, controlled for changes in the number of children	187
2.F.3	Event-study estimates of the impact of the childcare shock on the supply of individualized childcare, by treatment group, controlled for changes in the number of children	188
2.F.4	Event-study estimates of the impact of the childcare shock on paid parental leave take-up, by treatment group, with département-level calendar time fixed effects	190
2.F.5	Event-study estimates of the impact of the childcare shock on paid parental leave take-up, by treatment group, with Zone d'emploi-level calendar time fixed effects	191
2.F.6	Event-study estimates of the impact of the childcare shock on paid parental leave take-up, by treatment group, with Bassin de vie-level calendar time fixed effects	192

2.F.7	Event-study estimates of the impact of the childcare shock on the supply of individualized childcare, by treatment group, with département-level calendar time fixed effects	193
2.F.8	Event-study estimates of the impact of the childcare shock on the supply of individualized childcare, by treatment group, with Zone d'emploi-level calendar time fixed effects	194
2.F.9	Event-study estimates of the impact of the childcare shock on the supply of individualized childcare, by treatment group, with Bassin de vie-level calendar time fixed effects	195
3.1	Lifecycle profile of fertility: share of mothers among nurses of both genders	220
3.2	Lifecycle profile of hospital nurses' labor supply: total labor supply in the salaried sector	221
3.3	Lifecycle profile of hospital nurses' labor supply: decompositions	222
3.4	Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector	223
3.5	Event-study estimates of the impact of children on mothers' labor supply: decompositions	224
3.6	Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector	225
3.7	Contribution of children to the lifecycle profile of nurses' labor supply: decompositions	226
3.A.1	Distribution of hours worked for part-time workers, 1995-1998	231
3.B.1	Share of nurses who ever hold a job at a hospital	233
3.C.1	Lifecycle profile of hospital nurses' labor supply: decompositions at the extensive margin	234
3.C.2	Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector, by fertility margin	237
3.C.3	Event-study estimates of the impact of children on mothers' labor supply: decompositions at the extensive margin	238
3.C.4	Event-study estimates of the impact of children on mothers' labor outcomes: labor earnings	239

3.C.5	Event-study estimates of the impact of children on fathers' labor supply: total labor supply in the salaried sector	240
3.C.6	Event-study estimates of the impact of children on mothers' labor outcomes: hours worked including overtime and hourly wages	242
3.C.7	Contribution of children to the lifecycle profile of nurses' labor supply: decompositions at the extensive margin	243
3.D.1	Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector	245
3.D.2	Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector	246
3.D.3	Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector	248
3.D.4	Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector	249
3.D.5	Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector	251
3.D.6	Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector	252
3.D.7	Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector	254
3.D.8	Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector	255
4.1	Plant size before mass layoff	285
4.2	Mass layoff intensity	286
4.3	Difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, for those who were before displacement	287
4.4	Difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, for those who were not before displacement	288
4.5	Difference-in-difference estimates of the impact of job displacement on the number of children born to workers	289

4.6	Difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status	290
4.7	Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status	291
4.8	Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the plant level	292
4.9	Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the industry level	293
4.10	Difference-in-difference estimates of the impact of job displacement on workers' households' overall income	294
4.11	Plant size and outflows over time	295
4.C.1	Difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status	311
4.D.1	Difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status (workers in a cohabiting heterosexual relationship)	313
4.E.1	Difference-in-difference estimates of the impact of job displacement on workers' presence in the income tax returns data	314
4.E.2	Difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship (balanced panel)	315
4.E.3	Difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status (balanced panel)	316
4.E.4	Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status (balanced panel)	317
4.E.5	Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the plant level (balanced panel)	318

4.E.6	Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the industry level (balanced panel)	319
4.E.7	Difference-in-difference estimates of the impact of job displacement on workers' households' overall income (balanced panel)	320
4.E.8	Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status (balanced panel of spouses)	322
4.E.9	Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the plant level (balanced panel)	323
4.E.10	Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the industry level (balanced panel)	324
4.F.1	Difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, for those who were before displacement (plant-closure events only)	325
4.F.2	Difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, for those who were not before displacement (plant-closure events only) . .	326
4.F.3	Difference-in-difference estimates of the impact of job displacement on the number of children born to workers (plant-closure events only) . .	327
4.F.4	Difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status (plant-closure events only)	328
4.F.5	Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status (plant-closure events only)	329

4.F.6	Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the plant level (plant-closure events only)	330
4.F.7	Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the industry level (plant-closure events only)	331
4.F.8	Difference-in-difference estimates of the impact of job displacement on workers' households' overall income (plant-closure events only)	332

List of Tables

1.1	Sample selection	67
1.H.1	Correlation between childcare preferences and hourly wages	95
2.1	Data description	132
2.2	Summary statistics	133
2.3	Event-study estimates of the impact of childcare expansions on parents' labor outcomes	134
2.4	Instrumental variable estimates of the impact of affordable collective childcare on parents' labor outcomes, by gender	135
2.C.1	Summary statistics at the municipality level: by treatment group	155
2.D.1	Empirical policy evaluation: counterfactual scenarios	160
2.F.1	OLS estimates of the association between observable characteristics in the 2006 Census and the timing of the local childcare expansion	169
2.F.2	Instrumental variable estimates of the impact of affordable collective childcare on parents' labor outcomes, by gender	171
2.F.3	Instrumental variable estimates of the impact of affordable collective childcare on parents' labor outcomes, by gender	173
2.F.4	Instrumental variable estimates of the impact of affordable collective childcare on parents' labor outcomes, by gender	174
2.F.5	Instrumental variable estimates of the impact of affordable collective childcare on parents' labor outcomes, by gender	175
2.F.6	Instrumental variable estimates of the impact of affordable collective childcare on parents' labor outcomes, by gender	177
2.F.7	Instrumental variable estimates of the impact of affordable collective childcare on parents' labor outcomes based on the opening of the first EAJE-PSU facility, by gender	184

3.1	Detailed occupations distribution among selected jobs (2009-2017) . . .	227
3.2	Summary statistics	228
4.1	Summary statistics: occupation and family structure	296
4.2	Summary statistics: displaced workers' salaried earnings and labor supply two years before separation (payroll tax data)	297
4.3	Summary statistics: displaced workers' earnings and labor supply two years before separation (income tax returns)	298
4.4	Summary statistics: displaced workers' spouses' earnings and labor supply two years before separation (income tax returns)	299
4.5	Summary statistics: displaced workers' household income two years before separation (income tax returns)	300
4.F.1	Summary statistics: occupation and family structure	334
4.F.2	Summary statistics: displaced workers' salaried earnings and labor supply two years before separation (payroll tax data)	335
4.F.3	Summary statistics: displaced workers' earnings and labor supply two years before separation (income tax returns)	336
4.F.4	Summary statistics: displaced workers' spouses' earnings and labor supply two years before separation (income tax returns)	337
4.F.5	Summary statistics: displaced workers' household income two years before separation (income tax returns)	338

Introduction

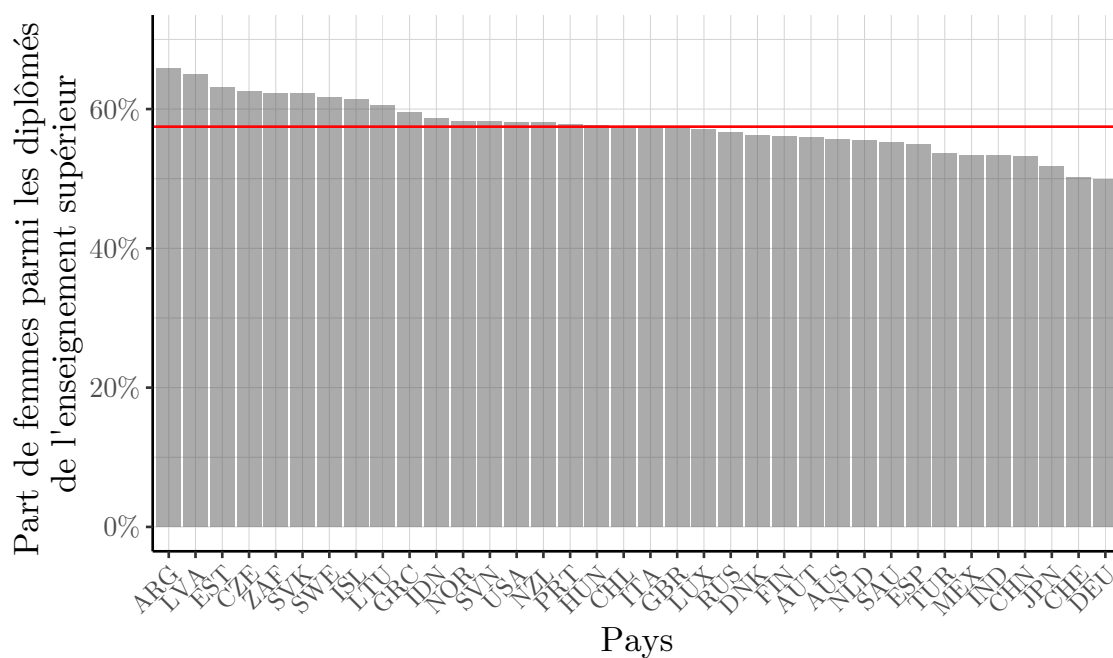
Vie familiale et écarts entre les sexes sur le marché du travail

Les inégalités entre les hommes et les femmes sur le marché du travail constituent l'un des faits stylisés les plus saillants et robustes, aussi bien dans les économies avancées que dans les pays en voie de développement : les femmes représentent une part plus faible de l'emploi, et lorsqu'elles sont salariées, travaillent moins d'heures et pour des salaires plus faibles que leurs homologues masculins. L'ampleur de ces écarts s'est toutefois considérablement réduite au cours de la seconde moitié du XXe siècle. Cette évolution favorable résulte en grande partie de la réduction des écarts de niveau d'éducation entre femmes et hommes, et de la participation plus continue de celles-ci au marché du travail. Comme l'expose [Goldin \(2006\)](#), cette « *révolution tranquille* » découle très largement de la maîtrise de la fécondité grâce à la diffusion de la pilule contraceptive dans les années soixante (voir également [Goldin et Katz, 2002](#)). Les jeunes femmes ont ainsi pu maîtriser le calendrier de leurs maternités, envisager des études longues et se projeter dans des carrières continues. Une partie des écarts passés entre les sexes résultait en effet de différences de productivité sur le marché du travail, liées à des études plus courtes et à une moindre expérience accumulée sur le marché du travail. Aujourd'hui, bien que femmes et hommes fassent des choix d'éducation différents en moyenne, par exemple en termes de filières universitaires ([Zafar, 2013](#)), les écarts entre les sexes en matière de niveau d'éducation se sont inversés dans presque tous les pays développés (voir [Figure 1](#)). Il semble par conséquent peu probable que les inégalités de genre qui persistent sur le marché du travail proviennent essentiellement d'écarts de productivité.

L'impact des maternités sur les inégalités entre femmes et hommes sur le marché du travail est néanmoins loin d'avoir disparu. Bien au contraire, des recherches récentes montrent que, dans les pays développés, la grande majorité des écarts qui subsistent entre les sexes sur le marché du travail résultent des conséquences *directes*¹ de la fécondité sur les carrières des mères ([Kleven et al., 2019b](#)). Ces effets négatifs étaient certes connus et analysés depuis longtemps par économistes et sociologues

¹Par opposition aux conséquences *anticipées* de la fécondité sur les décisions d'éducation discutée au paragraphe précédent.

Figure 1 – Part de femmes parmi les nouveaux diplômés de l’enseignement supérieur, 2019



Part des femmes parmi les nouveaux diplômés de l’enseignement supérieur, par pays. La barre rouge représente la moyenne pour les pays de l’OCDE. Le champ comprend les diplômés de l’enseignement supérieur court, ainsi que les diplômés de niveau licence et master. La période couverte est 2019 ou la dernière année disponible.

Source : OCDE, EIS et Eurostat.

(par exemple par [Waldfogel, 1995, 1997](#)), tout comme l'effet de la fécondité sur l'offre de travail des femmes (voir par exemple [Fleisher et Rhodes, 1979](#)). Cependant, ce courant relativement récent en économie du travail diffère de ces approches plus anciennes tant sur la méthode que sur la portée des résultats. D'une part, au contraire de la plupart des travaux antérieurs sur l'écart des rémunérations entre les femmes et les hommes qui utilisent les approches traditionnelles de décomposition ([Blinder, 1973](#); [Oaxaca, 1973](#)), il se fonde sur des techniques quasi-expérimentales pour identifier de manière crédible l'effet causal de la fécondité sur l'emploi et les carrières des parents. D'autre part, la cible des nouvelles approches est clairement définie comme l'effet causal de la maternité sur les trajectoires professionnelles des mères dans l'ensemble de la population, qui est la grandeur d'intérêt lorsqu'il s'agit de considérer les écarts agrégés entre les sexes. En cela, ces nouvelles approches contrastent avec certains travaux antérieurs : dans leur recherche d'instruments particulièrement convaincants pour éliminer les biais d'endogénéité et identifier les conséquences des décisions de fécondité sur l'emploi des mères, certains auteurs ont pu se concentrer sur des sous-populations assez sélectionnées, ce qui rendait l'extrapolation des effets estimés assez délicate (voir par exemple [Angrist et Evans, 1998](#); [Lundborg et al., 2017](#)).

Ces nouveaux développements scientifiques ont été rendus possibles par la disponibilité accrue de données administratives détaillées et complètes pour les chercheurs ([Card et al., 2010](#)). Cela explique pourquoi les pays d'Europe du Nord comme la Suède ([Angelov et al., 2016](#)) ou le Danemark ([Kleven et al., 2019b](#)), pour lesquels les versions statistiques des registres administratifs sont depuis longtemps accessibles aux chercheurs, ont été les premiers domaines d'application de ces méthodes. Des comparaisons internationales ont depuis été effectuées pour quelques pays ([Kleven et al., 2019a](#)). Dans le cas français, [Meurs et Pora \(2019\)](#) montrent que l'éducation et l'expérience n'expliquent plus l'écart de salaire entre les femmes et les hommes, et fournissent des estimations de l'effet de la maternité sur le revenu salarial des mères basées sur le cadre développé par [Kleven et al. \(2019b\)](#). En termes de comparaisons internationales, leur estimation de l'effet de long terme, 32% (voir [Figure 2](#)) se situe entre celles des pays scandinaves (Danemark, 21%; Suède, 26%) et celles des pays anglophones (Royaume-Uni, 44%; États-Unis, 31%). Cette perte de revenu sala-

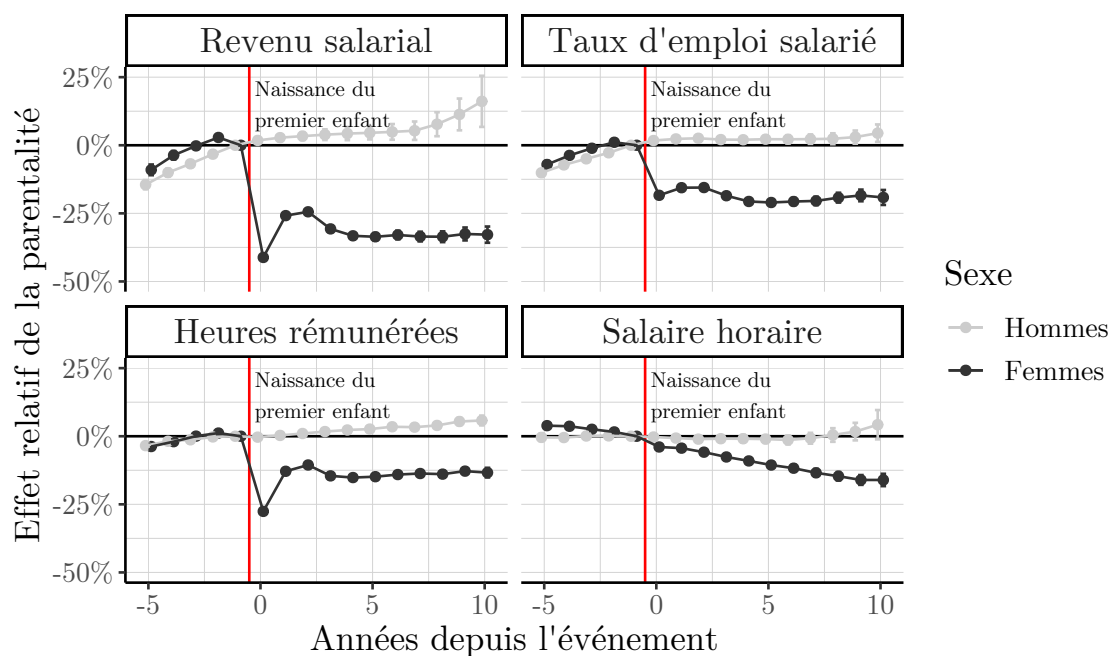
rial due à l'arrivée des enfants découle essentiellement d'une diminution du temps consacré au travail salarié, que ce soit par passage à temps partiel ou par retrait de l'activité salariée.

L'offre de travail des mères : quelles théories en concurrence ?

D'un point de vue théorique, ces résultats empiriques soulignent que les individus ne prennent pas leurs décisions d'allocation de temps et d'effort de manière isolée, comme le supposeraient les modèles d'offre de travail les plus simples. Au contraire, ils sont étroitement liés par un réseau de relations qui délimite des unités au sein desquelles la plupart des décisions sont prises conjointement : les familles. Ce fait est reconnu depuis longtemps par ce que certains ont appelé la nouvelle économie de la famille (Becker, 1974). En effet, lorsque les individus sont considérés de manière isolée, l'allocation de leur temps devrait seulement dépendre d'une comparaison de leurs propres productivités entre les secteurs, à savoir le rendement du temps passé sur le marché du travail par rapport au temps passé à la production domestique. En revanche, lorsqu'ils sont considérés comme faisant partie d'une famille dans laquelle, en dehors de différences de productivité, le temps et l'effort consacrés par chaque individu à la production domestique sont un substitut parfait du temps et de l'effort consacrés par tout autre membre de la famille, l'allocation du temps et de l'effort doit refléter la structure de l'avantage comparatif au sein de la famille (Becker, 1981). En conséquence, le fait qu'un individu soit plus productif sur le marché du travail que dans la production domestique ne suffit plus à expliquer qu'il consacre plus de temps au premier qu'au second : la décision dépend en fait de la productivité relative des autres membres de la cellule familiale.

Pris au pied de la lettre, le modèle de famille développé par Becker, statique, certain et unitaire, dans lequel les individus se substituent parfaitement les uns aux autres, est rejeté par les données. En effet, il prédit qu'au plus un membre du ménage consacrera du temps et des efforts à la fois au marché du travail et à la production domestique. En d'autres termes, pour rationaliser les différences entre femmes et

Figure 2 – Estimations de l'effet de la parentalité sur la position des parents sur le marché du travail, France 2002-2015



France métropolitaine, ensemble des salariés du secteur privé âgés de 20 à 59 ans, hors salariés agricoles, ayant travaillé au moins une heure dans le secteur privé entre 1995 et 2015, avec un nombre d'heures par jour supérieur à 1/8e de la durée légale du travail et un salaire horaire supérieur à 95% du Smic horaire.

Source : [Meurs et Pora \(2019\)](#).

hommes en matière d'allocation du temps, ce modèle prédirait que les hommes qui consacrent un temps non-nul à la production domestique doivent également être les seuls pourvoyeurs de revenu de leur famille, leurs conjointes se spécialisant entièrement dans la production domestique. Cette prédiction contredit la situation actuelle dans laquelle les couples bi-actifs sont les plus nombreux, les deux conjoints consacrant une part non-nulle quoique très différente de leur temps à la production domestique (Bart *et al.*, 2015).

En plus de cette invalidation empirique, cette approche a depuis été contestée sur des bases théoriques, parce qu'elle enfreint l'individualisme méthodologique : ce modèle *unitaire* de ménage ne peut refléter une décision prise par les différents individus qui le composent que dans la mesure où ceux-ci mettent toutes leurs ressources en commun, et ont les mêmes préférences (Chiappori, 1988). Une partie importante de l'économie de la famille a donc développé des modèles *collectifs* du ménage, en mettant l'accent sur la règle de décision au sein des familles et en incorporant progressivement des aspects importants de la vie familiale tels que la production domestique (Chiappori, 1997) et les enfants (Blundell *et al.*, 2005). Une limite de l'approche des modèles collectifs est que ce haut niveau de flexibilité rend difficile la mise à l'épreuve empirique des conclusions théoriques, faute de données. Par exemple, lorsqu'il incorpore la production domestique et les enfants, le modèle ne peut être identifié qu'avec des données détaillées sur l'utilisation du temps qui permettent de distinguer le temps consacré à la production domestique de celui dédié aux loisirs.

Dans le contexte particulier des différences entre femmes et hommes en matière de décisions relatives aux enfants, une difficulté supplémentaire a trait à l'importance des normes et de l'identité. En effet, l'arrivée des enfants est susceptible de modifier non seulement la structure de l'avantage comparatif au sein de la famille, mais aussi le jugement des parents quant à l'allocation appropriée de leur temps et de leur effort compte tenu de leur appartenance à l'un ou l'autre sexe. Cet effet englobe une multiplicité de canaux que les économistes ont généralement regroupés sous le terme d'identité. Théoriquement, l'identité de genre est parfois modélisée comme un coût psychologique que les femmes et les hommes subissent lorsque leurs

choix s'écartent de comportements définis de manière exogène et jugés appropriés pour leur genre (Akerlof et Kranton, 2000). Cependant, d'autres auteurs choisissent plutôt de considérer les écarts par rapport aux comportements médians à l'équilibre (Barigozzi *et al.*, 2018). Ce dernier choix est plus coûteux en termes de modélisation, car il est dépourvu de sens en dehors du contexte d'une population d'agents hétérogènes. En effet, en l'absence d'hétérogénéité, tous les ménages sont identiques et ne peuvent donc pas dévier du comportement du ménage médian. Cependant, il a l'avantage de correspondre à des phénomènes d'*effets de pairs* qui ont été documenté empiriquement (Maurin et Moschion, 2009). Le principal intérêt de cette approche est de surcroît de permettre d'appliquer les outils normatifs usuels de l'économie publique aux normes de genre. En effet, dans ce modèle, si elles sont assez nombreuses, les familles qui font le choix traditionnel de consacrer une part importante du temps des mères à l'éducation des enfants accroissent le coût pour les autres familles d'une allocation du temps dans laquelle les mères passent davantage de temps sur le marché du travail. En d'autres termes, l'équilibre de *laisser-faire* n'est pas efficace, parce que les décisions d'allocation du temps d'une famille imposent une externalité sur les autres. Cela justifie l'intervention du décideur public, qui peut par exemple décider de subventionner des services de garde d'enfant.

Structure et objet de la thèse

Cette thèse étudie de façon empirique l'effet de la vie familiale sur l'offre de travail dans le contexte français. Elle comprend quatre chapitres indépendants, chacun consacré à un aspect précis de l'impact de la vie familiale sur les décisions d'allocation de temps au travail à l'extérieur du foyer. Le chapitre 1, co-écrit avec Lionel Wilner, contribue à la littérature récente traitant de l'effet de la fécondité sur les revenus et l'offre de travail des mères en examinant l'hétérogénéité de ces effets le long de la distribution de leurs salaires potentiels. Le chapitre 2 cherche à déterminer si l'augmentation de la disponibilité de services de garde d'enfants abordables fournis par des crèches subventionnées par les pouvoirs publics, et donc la diminution du coût implicite du travail à l'extérieur du domicile, permet aux parents, et en particulier aux mères de jeunes enfants, d'augmenter leur offre de travail et leurs revenus. Le chapitre 3 se concentre sur les conséquences de la vie familiale sur les systèmes

de santé, en quantifiant la contribution de la parentalité à la diminution de l'offre de travail des infirmières au long de leurs carrières. Enfin, le chapitre 4, co-écrit avec Raphaël Lardeux, explore comment un changement brutal et permanent, ici le licenciement d'un des adultes de la famille, affecte la structure de l'avantage comparatif intra-familial et se répercute sur les relations familiales et l'offre de travail des autres membres.

Données et méthodes

Ces chapitres sont liés non seulement par leur objet commun, l'effet de la vie familiale sur l'offre de travail, mais aussi par les données et les méthodes qu'ils mobilisent. Tout d'abord, les quatre chapitres reposent sur des données issues de registres administratifs français appelés DADS (*Déclarations Annuelles de Données Sociales*) et EDP (*Échantillon Démographique Permanent*). Les DADS regroupent les informations fournies par les employeurs lors du remplissage de formulaires de déclaration de salaires : revenus salariaux, heures et jours travaillés, profession, secteur d'activité, ainsi que des caractéristiques individuelles telles que l'âge et le sexe. Les DADS ont l'avantage de permettre le rapprochement entre employeur et salarié car elles intègrent le Siret (*Système d'Identification au Répertoire des Établissements*), un identifiant d'établissement. Ce lien salarié-employeur est crucial pour le dernier chapitre de cette thèse, qui étudie les conséquences sur les décisions prises par les salariés et leurs familles de chocs qui se produisent au niveau de l'établissement où ils travaillent.

Les données de l'EDP combinent des informations administratives provenant de sources multiples, les plus utiles ici étant les bulletins de naissance et les déclarations fiscales pour l'impôt sur le revenu et la taxe d'habitation. Les données regroupées dans les DADS et l'EDP peuvent être appariées grâce à un identifiant anonyme basée sur le NIR (*Numéro d'Inscription au Répertoire*), un numéro de sécurité sociale. Cette possibilité est particulièrement importante pour cette thèse car elle permet de savoir si un individu adulte a des enfants, et les dates de naissance de ces enfants. La limite la plus importante de ces données est cependant qu'elles ne couvrent qu'un

échantillon de la population française², défini au niveau individuel, par opposition au niveau du ménage, ce qui limite la possibilité de considérer les couples de manière symétrique.

Les quatre chapitres s'appuient sur la nature longitudinale des données et sur l'hétérogénéité de l'exposition aux divers chocs et ont donc recours à une variante de la technique canonique de différence-de-différences pour identifier l'effet causal de la vie familiale sur l'offre de travail. Initiée en épidémiologie par [Snow \(1855\)](#), cette approche est devenue commune chez les économistes du travail et les spécialistes des finances publiques au cours des dernières décennies (voir par exemple [Card et Krueger, 1994](#)), dans la lignée de ce que certains ont pu appeler « *révolution de la crédibilité* » ([Angrist et Pischke, 2010](#)). La thèse met en œuvre ce cadre général à différentes échelles. Les chapitres 1 et 3 considèrent l'exposition à la parentalité, qui correspond aux naissances de ses enfants et est donc définie au niveau individuel ou familial. Le chapitre 2 se concentre sur les chocs définis par des augmentations soudaines de la disponibilité des places en crèche, mesurée au niveau de la commune. Enfin, le chapitre 4 étudie les cessations d'emploi exogènes déclenchées par des événements qui ont lieu au niveau de l'établissement. Des économètres ont récemment mis en doute le bien fondé de variations courantes de cette stratégie pour identifier les effets causaux d'intérêt, soit lorsque la différence de différences est utilisée comme un instrument ([de Chaisemartin et D'Haultfoeulle, 2017](#)), soit lorsque l'exposition au traitement d'intérêt a lieu à des dates qui varient d'une unité à l'autre ([Callaway et Sant'Anna, 2020](#) ; [de Chaisemartin et D'Haultfoeulle, 2020](#) ; [Sun et Abraham, 2020](#) ; [Goodman-Bacon, 2021](#)). Les quatre chapitres de cette thèse relèvent tous d'au moins une de ces catégories et intègrent donc des méthodes issues de cette littérature pour remédier à ces difficultés.

Résumé des chapitres et contributions

Chap. 1 : « Dissecting Child Penalties » co-écrit avec Lionel Wilner

Ce chapitre vise à comprendre les mécanismes qui produisent l'effet négatif des enfants sur les carrières des mères. En effet, celui-ci explique désormais la plus grande


²Le taux d'échantillonnage est de 4,4% pour la période la plus récente.

partie des écarts restants entre les sexes sur le marché du travail (Kleven *et al.*, 2019b). Le versant biologique de la maternité proprement dite et les contraintes de santé qui y sont liées n'expliquent que très peu la baisse des revenus des mères (Kleven *et al.*, 2021). De même, les productivités relatives du marché du travail, résumées par les revenus potentiels relatifs, ne jouent que très peu. L'avantage comparatif à la Becker (1981) ne semble donc pas l'explication la plus crédible. En effet, il faudrait pour cela que les mères aient une productivité bien plus élevée que les pères dans les activités liées aux soins aux enfants, sans que cette différence de productivité ne soit liée à des différences biologiques.

Dans ce chapitre, nous mettons l'accent sur la productivité absolue des mères sur le marché du travail, par opposition à leur productivité relative au sein du couple, en tant que déterminant clé de l'effet des enfants sur les carrières de leurs mères. Pour ce faire, nous développons un cadre qui intègre une version non-linéaire de différence-de-différences (*event-study*) en opérant un classement non-paramétrique des individus selon leur salaire horaire mesuré sur plusieurs années et avant la naissance de l'enfant. Comme les salaires horaires (potentiels) d'un même individu sont fortement corrélés dans le temps, ce classement approxime le classement des mères en fonction de leur salaire horaire potentiel (éventuellement inobservé) après la naissance de l'enfant. Notre approche de différence-de-différences compare les mères qui viennent d'avoir leur premier (n ième) enfant avec les femmes qui restent sans enfant (qui choisissent d'avoir exactement $n - 1$ enfants). Elle s'appuie sur une version légèrement modifiée et non-linéaire de l'hypothèse habituelle de tendances parallèles conditionnelles. Un avantage de cette modification est qu'elle est compatible avec une décomposition comptable de l'effet de la fécondité sur le revenu salarial en une somme d'ajustements des différentes composantes de l'offre de travail d'une part et de variation de salaire horaire d'autre part.

Nous appliquons ce cadre à une combinaison de données françaises de salaires et de bulletins de naissances couvrant la période 1995-2015 (Figure 3). Nos principaux résultats sont les suivants : (i) les femmes dont les salaires sont les plus élevés subissent des pertes de revenus du travail beaucoup plus faibles en raison de la naissance d'un enfant que celles moins bien rémunérées ; (ii) elles sont beaucoup moins susceptibles d'interrompre leur carrière ou de réduire leurs heures rémunérées ; (iii) fait important, les ampleurs de ces deux derniers effets présentent un compor-

Figure 3 – Estimations de l’effet de la parentalité sur les écarts de rémunération et d’offre de travail entre femmes et hommes, par position initiale dans la distribution de salaire



Estimations en triples différences de l’effet de la naissance du premier enfant sur l’écart de revenu salarial entre femmes et hommes et ses composantes. La date de référence est prise l’année qui précède immédiatement la naissance (contrefactuelle) du premier enfant. France métropolitaine, ensemble des salariés du secteur privé âgés de 20 à 59 ans, hors salariés agricoles. L’emploi salarié est restreint aux postes avec un nombre d’heures par jour supérieur à 1/8e de la durée légale du travail et un salaire horaire supérieur à 95% du Smic horaire.

Source : Insee, panel tous salariés et EDP.

tement monotone le long de la distribution des salaires horaires ; (iv) les pertes de salaire horaire semblent plutôt homogènes le long de la distribution des salaires horaires avant la naissance des enfants ; (v) l'hétérogénéité de l'effet des enfants sur les carrières des mères le long de la distribution des salaires horaires persiste à long terme, jusqu'à 10 ans après la naissance des enfants. En revanche, les hommes ne voient pas leurs revenus salariaux diminuer et ne réduisent pas leur offre de travail en réponse à l'arrivée des enfants, quelle que soit leur position dans la distribution des salaires. Nous fournissons des éléments supplémentaires qui montrent que ces résultats ne sont pas dus à un retour à la moyenne, à une sélection différenciée dans la parentalité ou à des différences de préférences qui entraîneraient des trajectoires d'accumulation de capital humain divergentes.

Ces résultats sont compatibles avec des modèles dans lesquels les familles arbitrent entre le temps que les mères passent sur le marché du travail et le temps qu'elles consacrent aux enfants lorsqu'elles prennent leur décision d'offre de travail, mais ne tiennent pas compte de cet arbitrage lorsqu'il s'agit des pères. En d'autres termes, les familles semblent considérer les services marchands comme des substituts acceptables des soins aux enfants fournis par les mères : leurs décisions paraissent refléter une comparaison entre les revenus générés par le temps passés par les mères sur le marché du travail, avec les coûts non-monétaires induits par ce temps, notamment pour la garde d'enfant. En revanche, elles ne semblent pas percevoir les contributions paternelles de la même manière. Cela expliquerait pourquoi, dans la littérature existante, les mères de jeunes enfants modifient leur comportement en réponse aux réformes qui rendent l'interruption de carrière liée aux enfants plus ou moins coûteuse (Piketty, 2005 ; Lequien, 2012 ; Joseph *et al.*, 2013), alors que les réformes récentes qui visent spécifiquement les pères n'ont pratiquement aucun impact sur leurs comportements (Périver et Verdugo, 2021).

Chap. 2 : « Keep Working and Spend Less? Collective Childcare and Parental Earnings in France »

Ce chapitre vise à évaluer l'efficacité des politiques françaises d'accueil de la petite enfance du point de vue de l'offre de travail des parents, et surtout des mères. Contrairement à d'autres pays, les dispositifs de garde d'enfants en France sont extrêmement diversifiés : une longue histoire institutionnelle a conduit à la coexistence

de congés parentaux rémunérés et de services de garde d'enfants formels fortement subventionnés, ces derniers englobant un continuum qui va de la garde individuelle à domicile aux services collectifs fournis par les crèches (Thévenon, 2011). Cette variété des modes de garde est encouragée par les décideurs politiques et le grand public, et est supposée offrir aux familles la liberté de choisir les modalités de garde d'enfants les plus adaptées à leurs différentes préférences et contraintes.

Cette diversité n'est pas sans conséquence : elle laisse place à des effets de substitution potentiellement importants entre les solutions de garde d'enfants. Pour tester cette hypothèse, j'étudie les conséquences d'une série de plans nationaux lancés dans les années 2000 et destinés à augmenter l'offre de crèches offrant un accueil collectif des très jeunes enfants particulièrement accessible du point de vue du prix pratiqué. Plus précisément, je m'intéresse aux répercussions de ces plans en termes de revenus du travail et d'offre de travail des parents et en termes de choix de modes de garde.

En m'appuyant sur un cadre similaire à celui de Duflo (2001), je mets à profit des *variations dans le calendrier des expansions de crèches entre communes*, dans des groupes de communes *qui ont connu des expansions d'ampleur similaire* (Figure 4), pour identifier l'effet causal des crèches fortement subventionnées sur les revenus du travail et l'offre de travail des parents, ainsi que sur leurs choix en matière de garde d'enfants. J'applique ce cadre à une série de données administratives détaillées : des registres de crèches et de congés parentaux tenus par la branche Famille de la Sécurité sociale française, ainsi que des bulletins de naissances et des données de salaires issues de l'EDP et des DADS.

Je constate que ces fortes augmentations de l'offre de garde collective à prix modéré au niveau des communes n'ont pas entraîné de changement substantiel dans les trajectoires professionnelles des parents. Plus précisément, mes estimations sont incompatibles avec des effets causaux sur l'emploi maternel supérieurs à 0,05 point de pourcentage par point de pourcentage d'augmentation du taux de couverture des services de garde d'enfants. Pour interpréter ces effets nuls, je reproduis mon analyse au niveau de la commune, en me concentrant cette fois sur l'offre d'autres solutions de garde d'enfants formelles et plus coûteuses, à savoir les assistantes maternelles et les nourrices qui gardent les enfants à domicile. Je mets en évidence un effet d'éviction très importante des crèches collectives sur ces solutions plus individualisées de

garde d'enfants. Dans les communes où l'offre de garde collective abordable a le plus fortement augmenté, la baisse à moyen terme de l'offre de garde individuelle est d'une ampleur équivalente à celle de l'augmentation de l'offre collective. Cela implique que l'augmentation de la capacité d'accueil des crèches profite probablement à des parents qui se seraient autrement tournés vers des solutions de garde formelles individualisées et plus coûteuses.

Chap. 3 : « Do Children Explain Nurses Shortages ? »

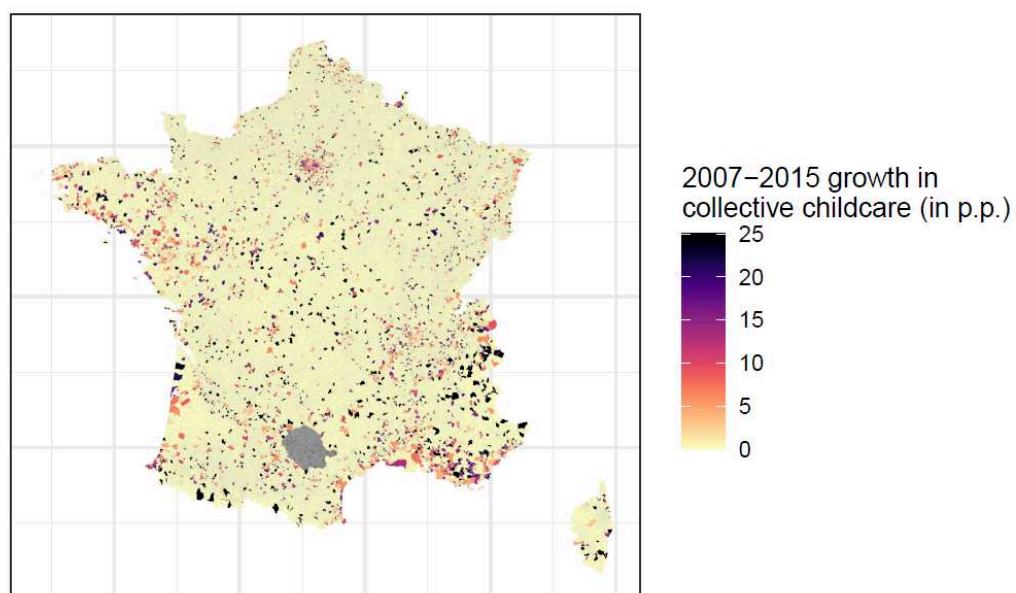
Ce chapitre quantifie la contribution de l'effet causal de la maternité sur l'offre de travail des infirmières hospitalières françaises au profil agrégé d'offre de travail de cette profession. En effet, l'insuffisance de l'offre d'infirmières est une préoccupation dans la plupart des économies avancées depuis plus de deux décennies (voir par exemple [Shields, 2004](#)). Le travail infirmier ayant un impact direct et positif sur la santé des patients ([Propper et Reenen, 2010](#) ; [Gruber et Kleiner, 2012](#)), cette pénurie peut conduire à des effets négatifs sur la santé. Le caractère extrêmement féminisé de la profession d'infirmière et des travaux récents montrant que l'offre de soins infirmiers peut réagir vivement aux politiques de congé parental ([Friedrich et Hackmann, 2021](#)), suggèrent qu'une mauvaise conciliation entre vie familiale et vie professionnelle, ainsi que les normes de genre qui assignent aux mères la responsabilité de l'éducation des enfants, pourraient expliquer une part substantielle de l'insuffisance de l'offre de travail des infirmières.

Je m'appuie sur des données administratives longitudinales détaillées issues à la fois des déclarations de salaires et des déclarations de naissance pour explorer cette question. Tout d'abord, je montre que (i) l'offre de travail moyenne des infirmières hospitalières dans le versant salarié de la profession³ diminue considérablement dans les années qui suivent leur premier emploi à l'hôpital, et (ii) le nombre de mères parmi elles augmente très fortement au cours de la même période.

Je mets ensuite en œuvre un cadre de différence-de-différences (*event-study*) afin d'identifier l'effet causal des enfants sur l'offre de travail des mères. La maternité conduit les infirmières à diminuer leur offre de travail dans le secteur salarié d'environ 0,15 équivalent temps plein au cours des dix premières années suivant la naissance de

³Par opposition à l'exercice libéral de la profession infirmière.

Figure 4 – Répartition spatiale de l'évolution entre 2007-2015 de l'offre relative d'accueil collectif de jeunes enfants



Croissance 2007-2015 du taux de couverture des Établissements d'Accueil du Jeune Enfant financés par la Prestation de Service Unique, rapporté à la population d'enfants de 2 ans ou moins, par commune. Note : Les données concernant le département du Tarn sont omises.

Source. Cnaf, registre des EAJE-PSU ; Insee, bulletins de naissance.

leur premier enfant (Figure 5). Cette diminution est entièrement due à des transitions vers des postes à temps partiel, par opposition à des démissions de l'hôpital.

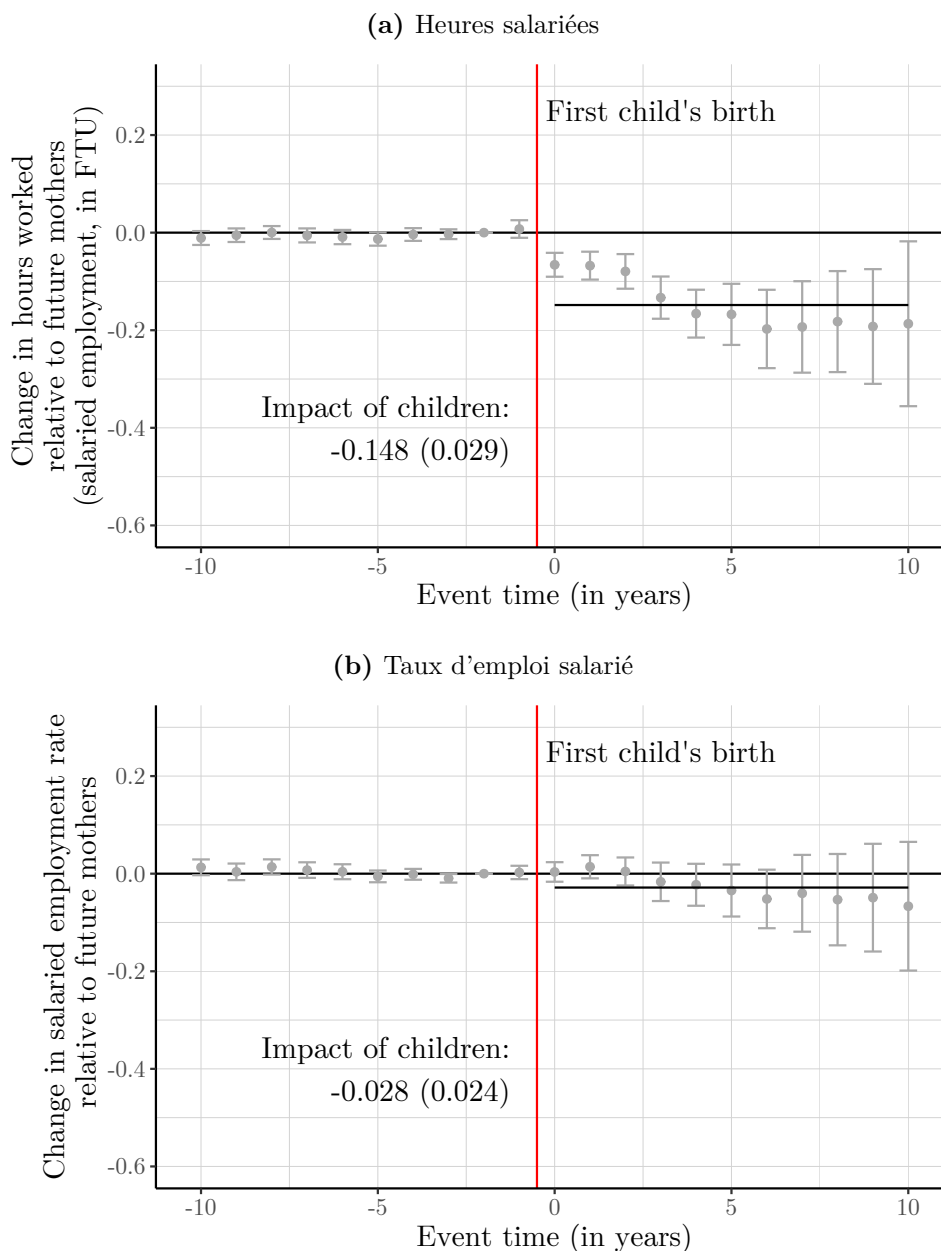
J'utilise enfin ces effets estimés pour comparer les profils de cycle de vie observés avec les profils contrefactuels qui seraient observés si l'effet des enfants était fixé à 0, c'est-à-dire (i) si les infirmières n'avaient pas d'enfants ou (ii) si les mères prenaient leurs décisions d'offre de travail de la même manière que les pères, puisque les pères ne réduisent pas leurs heures de travail à l'arrivée des enfants. Dans un tel cas, la diminution des heures travaillées en emploi salarié au cours des dix premières années de la carrière serait 37% moins importante, et serait réduite de moitié en ce qui concerne les heures travaillées dans le secteur public. Il y aurait en revanche toujours un nombre important d'infirmières quittant l'emploi salarié à l'hôpital : les maternités n'expliquent que très partiellement la diminution constatée de l'offre de travail.

Chap. 4 : « Job Displacement, Families and Redistribution » co-écrit avec Raphaël Lardeux

Ce chapitre étudie les conséquences d'une perte d'emploi sur de multiples dimensions de la structure familiale, du revenu individuel et du ménage, et des trajectoires sur le marché du travail. La croissance économique, qui exige la réallocation des travailleurs entre différents emplois (Aghion et Howitt, 1994), expose les salariés au risque d'être licenciés. Les divers systèmes de sécurité et d'assurance sociale qui caractérisent les économies modernes ont cherché à réduire les conséquences de ces transitions sur le bien-être des salariés, sans lesquelles l'effet du progrès technique serait sans équivoque positif (Aghion *et al.*, 2016). Cependant, ces systèmes d'assurance collective formels ne sont pas le seul canal par lequel les travailleurs peuvent atténuer ce risque. En effet, les familles peuvent offrir une bonne assurance (i) en diluant le risque individuel grâce à la mise en commun des ressources, et donc des chocs individuels de tous les membres du ménage⁴, et (ii) en modifiant leurs allocations de temps et d'effort pour s'adapter à la nouvelle structure de l'avantage comparatif au sein de la famille. Ce deuxième mécanisme est parfois connu sous le nom d'effet de travailleur supplémentaire (*added worker effect*, Lundberg, 1985).

⁴Cette mise en commun ne peut avoir de valeur d'assurance qu'à condition que les chocs ne soient pas trop fortement corrélés au sein de la famille.

Figure 5 – Estimations de l'effet de la fécondité sur l'offre de travail des infirmières hospitalières : offre de travail dans le secteur salarié



Estimations de l'impact des enfants sur les heures travaillées par les mères dans le secteur salarié, en équivalent temps plein, et sur le taux d'emploi salarié, selon le temps écoulé depuis la naissance du premier enfant. Les heures travaillées ne sont pas conditionnelles à l'emploi salarié, mais incorporent la marge de participation (0 heure travaillée). Les écarts types sont ajustés au niveau de clusters individuels et estimés par bootstrap avec 200 répliques.

Source. Insee, panel tous salariés et EDP.

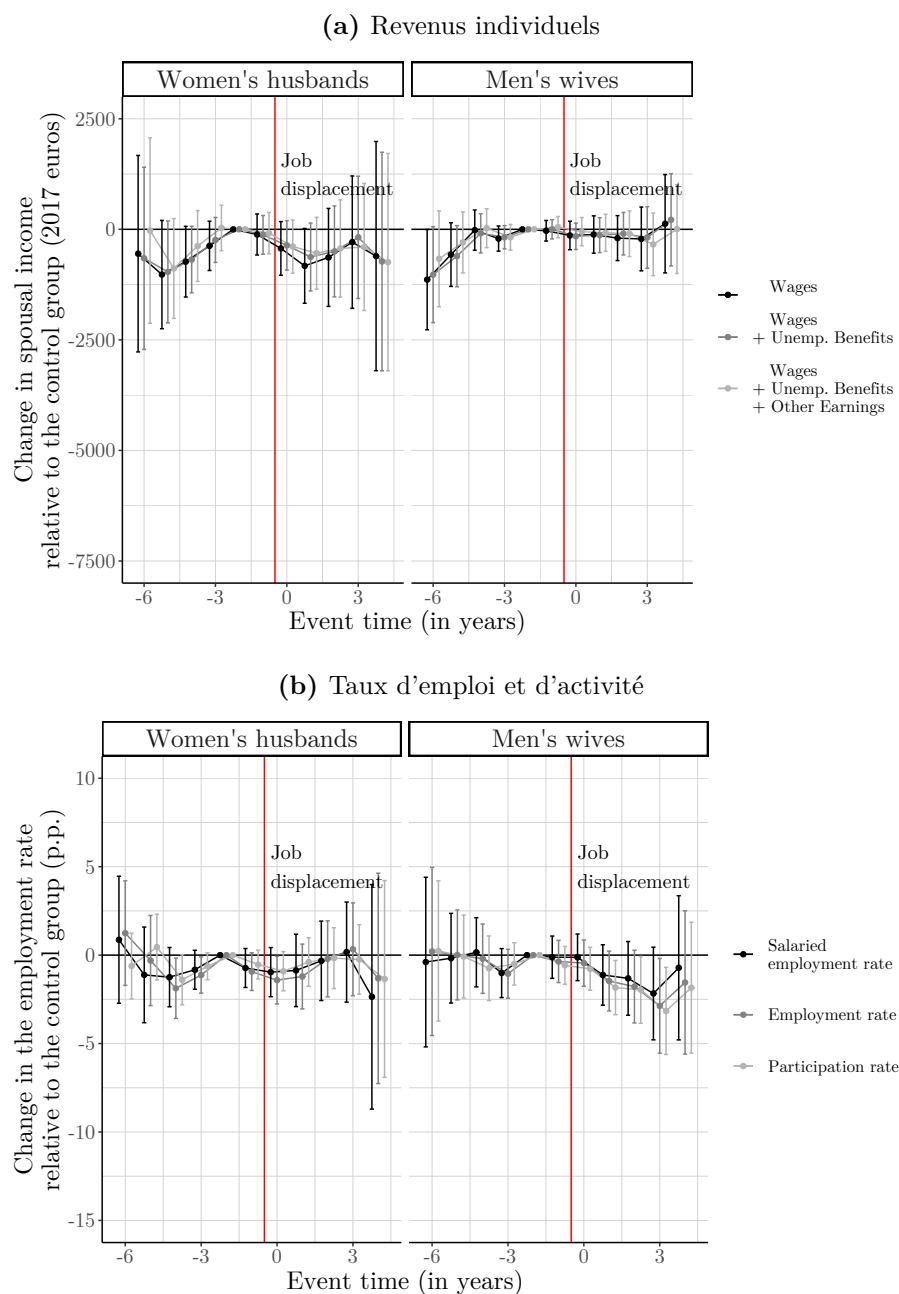
Cependant, si, pour le salarié licencié, les mécanismes d'assurance familiale augmentent la valeur de l'appartenance à une famille, le contraire est vrai des autres membres de la famille. Ces effets négatifs peuvent encore être amplifiés si les couples subissent un coût psychologique lorsque la nouvelle répartition du temps et des efforts est perçue comme contraire aux normes de genre traditionnelles. Compte tenu des rôles habituels des hommes et des femmes dans les pays occidentaux, cela peut être particulièrement le cas lorsque la perte d'emploi frappe les hommes.

En nous appuyant sur la méthodologie développée par [Jacobson *et al.* \(1993\)](#) et étendue par [Halla *et al.* \(2020\)](#), nous construisons un échantillon de travailleurs licenciés qui sont affectés par une perte d'emploi vraisemblablement exogène⁵, et nous les comparons à des travailleurs très similaires qui ne sont pas affectés par un tel choc. Ces comparaisons sont effectuées notamment à partir des données issues des déclarations d'impôts sur le revenu et de taxe d'habitation, ainsi que des bulletins de naissance.

Le licenciement n'a pratiquement aucun effet sur la structure familiale : la probabilité de rester dans une relation hétérosexuelle cohabitante pour les salariés qui étaient en couple avant d'être licenciés, ou la probabilité de mise en couple pour les célibataires ne changent pas en réponse à la perte d'emploi. De même, les décisions des salariés en matière de fécondité ne semblent pas réagir au choc. Nous ne trouvons pas d'indice d'un effet de travailleur supplémentaire : les revenus et l'offre de travail des conjoints des travailleurs licenciés n'augmentent pas en réponse à la baisse soudaine des revenus de leurs partenaires (Figure 6). En d'autres termes, si les familles fournissent une assurance contre le risque de revenu qui résulte du licenciement, c'est uniquement par la mise en commun des ressources plutôt que par des ajustements collectifs dans l'allocation du temps et des efforts. Enfin, alors que

⁵Le caractère plausiblement exogène de cette perte d'emploi tient au fait qu'elle découle d'un choc au niveau des établissements dans lequel travaillent les salariés licenciés, par opposition à une décision individuelle de quitter son emploi. En effet, nous exploitons la structure salarié-employeur des DADS pour identifier les établissements qui perdent brutalement 25% au moins de leur main d'œuvre entre deux années, et nous restreignons aux pertes d'emploi induites par ces brusques réductions de taille d'établissement.

Figure 6 – Estimations de l'effet de la perte d'emploi d'un salarié sur les revenus et l'offre de travail de son conjoint



Estimations par différence-de-différences repondérées de l'effet de la perte d'emploi d'un salarié sur l'offre de travail de son conjoint, par sexe et temps écoulé depuis le licenciement. Ces estimations sont corrigées de la corrélation entre pertes d'emploi à l'intérieur du couple. Les écarts-types sont ajustés au niveau de clusters définis par l'établissement où les salariés ont perdu leur emploi, et estimés par bootstrap avec 200 replications.

Source. Insee, DADS grand format ; Insee et DGFIP, EDP.

l'assurance chômage semble assurer les travailleurs de manière assez efficace contre le choc d'emploi de courte durée, elle perd de l'importance à long terme, la plupart des travailleurs réussissant à trouver un nouvel emploi. Les travailleurs apparaissent au mieux très partiellement assurés contre la composante de long terme du choc, grâce à un ensemble complexe qui comprend le travail indépendant, les marchés de capitaux et la progressivité du système d'impôts et de transferts. Cependant l'ampleur des pertes de revenu à long terme à la suite ce choc reste considérable, même après le retour à l'emploi et une fois prises en compte ces ressources.

Bibliographie

- AGHION, P., AKCIGIT, U., DEATON, A. et ROULET, A. (2016). Creative destruction and subjective well-being. *American Economic Review*, 106(12) :3869–97.
- AGHION, P. et HOWITT, P. (1994). Growth and unemployment. *Review of Economic Studies*, 61(3) :477–494.
- AKERLOF, G. A. et KRANTON, R. E. (2000). Economics and identity. *Quarterly Journal of Economics*, 115(3) :715–753.
- ANGELOV, N., JOHANSSON, P. et LINDAHL, E. (2016). Parenthood and the gender gap in pay. *Journal of Labor Economics*, 34(3) :545–579.
- ANGRIST, J. D. et EVANS, W. N. (1998). Children and their parents' labor supply : Evidence from exogenous variation in family size. *American Economic Review*, 88(3) :450–477.
- ANGRIST, J. D. et PISCHKE, J.-S. (2010). The credibility revolution in empirical economics : How better research design is taking the con out of econometrics. *Journal of Economic Perspectives*, 24(2) :3–30.
- BARIGOZZI, F., CREMER, H. et ROEDER, K. (2018). Women's career choices, social norms and child care policies. *Journal of Public Economics*, 168 :162–173.
- BART, V., CASTELL, L., ALLÈGRE, G., MARTIN, H. et LIPPMANN, Q. (2015). Travail domestique : les couples mono-actifs en font-ils vraiment plus ? *Economie et statistique*, (478-480) :189–208.
- BECKER, G. (1981). *A Treatise on the Family*. Harvard University Press, Cambridge.
- BECKER, G. S. (1974). On the relevance of the new economics of the family. *American Economic Review*, 64(2) :317–319.
- BLINDER, A. S. (1973). Wage discrimination : Reduced form and structural estimates. *Journal of Human Resources*, 8(4) :436–455.
- BLUNDELL, R., CHIAPPORI, P. et MEGHIR, C. (2005). Collective labor supply with children. *Journal of Political Economy*, 113(6) :1277–1306.
- CALLAWAY, B. et SANTANNA, P. H. (2020). Difference-in-differences with multiple time periods. *Journal of Econometrics*.
- CARD, D., CHETTY, R., FELDSTEIN, M. S. et SAEZ, E. (2010). Expanding access to administrative data for research in the United States. American Economic Association, ten years and beyond : Economists answer NSF's call for long-term research agendas.
- CARD, D. et KRUEGER, A. B. (1994). Minimum Wages and Employment : A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania. *American Economic Review*, 84(4) :772–793.
- CHIAPPORI, P.-A. (1988). Rational household labor supply. *Econometrica*, 56(1) : 63–90.
- CHIAPPORI, P.-A. (1997). Introducing household production in collective models of labor supply. *Journal of Political Economy*, 105(1) :191–209.
- de CHAISEMARTIN, C. et D'HAULTFŒUILLE, X. (2017). Fuzzy Differences-in-Differences. *Review of Economic Studies*, 85(2) :999–1028.

- de CHAISEMARTIN, C. et D’HAULTFŒUILLE, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9) :2964–96.
- DUFLO, E. (2001). Schooling and labor market consequences of school construction in indonesia : Evidence from an unusual policy experiment. *American Economic Review*, 91(4) :795–813.
- FLEISHER, B. M. et RHODES, G. F. (1979). Fertility, women’s wage rates, and labor supply. *American Economic Review*, 69(1) :14–24.
- FRIEDRICH, B. U. et HACKMANN, M. B. (2021). The Returns to Nursing : Evidence from a Parental-Leave Program. *Review of Economic Studies*.
- GOLDIN, C. (2006). The quiet revolution that transformed women’s employment, education, and family. *American Economic Review*, 96(2) :1–21.
- GOLDIN, C. et KATZ, L. F. (2002). The power of the pill : Oral contraceptives and womens career and marriage decisions. *Journal of Political Economy*, 110(4) :730–770.
- GOODMAN-BACON, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*.
- GRUBER, J. et KLEINER, S. A. (2012). Do strikes kill? evidence from new york state. *American Economic Journal : Economic Policy*, 4(1) :127–57.
- HALLA, M., SCHMIEDER, J. et WEBER, A. (2020). Job displacement, family dynamics, and spousal labor supply. *American Economic Journal : Applied Economics*, 12(4) :253–87.
- JACOBSON, L. S., LALONDE, R. J. et SULLIVAN, D. G. (1993). Earnings losses of displaced workers. *American Economic Review*, 83(4) :685–709.
- JOSEPH, O., PAILHÉ, A., RECOTILLET, I. et SOLAZ, A. (2013). The economic impact of taking short parental leave : Evaluation of a french reform. *Labour Economics*, 25 :63 – 75.
- KLEVEN, H., LANDAIS, C., POSCH, J., STEINHAEUER, A. et ZWEIMÜLLER, J. (2019a). Child penalties across countries : Evidence and explanations. *AEA Papers and Proceedings*, 109 :122–26.
- KLEVEN, H., LANDAIS, C. et SØGAARD, J. E. (2019b). Children and gender inequality : Evidence from Denmark. *American Economic Journal : Applied Economics*, 11(4) :181–209.
- KLEVEN, H., LANDAIS, C. et SØGAARD, J. E. (2021). Does biology drive child penalties? evidence from biological and adoptive families. *American Economic Review : Insights*, 3(2) :183–98.
- LEQUIEN, L. (2012). The impact of parental leave duration on later wages. *Annals of Economics and Statistics*, (107/108) :267–285.
- LUNDBERG, S. (1985). The added worker effect. *Journal of Labor Economics*, 3(1, Part 1) :11–37.
- LUNDBORG, P., PLUG, E. et RASMUSSEN, A. W. (2017). Can women have children and a career? IV evidence from IVF treatments. *American Economic Review*, 107(6) :1611–37.

- MAURIN, E. et MOSCHION, J. (2009). The social multiplier and labor market participation of mothers. *American Economic Journal : Applied Economics*, 1(1) :251–72.
- MEURS, D. et PORA, P. (2019). Gender equality on the labour market in France : A slow convergence hampered by motherhood. *Economie et Statistique / Economics and Statistics*, (510-511-5) :109–130.
- OAXACA, R. (1973). Male-female wage differentials in urban labor markets. *International Economic Review*, 14(3) :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*, (156) :79–109.
- PROPPER, C. et REENEN, J. V. (2010). Can pay regulation kill? panel data evidence on the effect of labor markets on hospital performance. *Journal of Political Economy*, 118(2) :222–273.
- PÉRIVIER, H. et VERDUGO, G. (2021). Can parental leave be shared? OFCE Working Paper 6, OFCE.
- SHIELDS, M. A. (2004). Addressing nurse shortages : what can policy makers learn from the econometric evidence on nurse labour supply? *Economic Journal*, 114(499) :F464–F498.
- SNOW, J. (1855). *On the mode of communication of cholera*. John Churchill.
- SUN, L. et ABRAHAM, S. (2020). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*.
- THÉVENON, O. (2011). Family policies in OECD countries : A comparative analysis. *Population and Development Review*, 37(1) :57–87.
- WALDFOGEL, J. (1995). The price of motherhood : Family status and women’s pay in young British cohort. *Oxford Economic Papers*, 47(4) :584–610.
- WALDFOGEL, J. (1997). The effect of children on women’s wages. *American Sociological Review*, 62(2) :209–217.
- ZAFAR, B. (2013). College major choice and the gender gap. *Journal of Human Resources*, 48(3) :545–595.

Chapter 1

Dissecting child penalties

joint with Lionel Wilner

1.1 Introduction

Recent research has highlighted that women’s earnings losses due to motherhood, referred to as child penalties, have become the main driver of gender inequality in the labor market in developed countries (Juhn and McCue, 2017; Kleven et al., 2019b). Surprisingly, biology explains actually very little of these penalties (Kleven et al., 2021). Similarly, within-couple relative labor market productivities, as measured by predicted potential earnings do not seem to explain why mothers are the ones who decrease their time and effort allocated to the labor market after the arrival of children. As a result, comparative advantage *à la* Becker (1981) is not the most compelling mechanism. Indeed, for this explanation to hold would require women to be much more productive than men in child-rearing activities, in ways totally unrelated to biology.

In this paper, we emphasize women’s *absolute* labor market productivity, as opposed to their *relative* within-household productivity, as a key determinant of their child-related labor supply decisions. Specifically, we show that the trade-off that mothers face between time spent outside the labor force, presumably devoted to child-rearing activities and home production, and their foregone labor earnings is a crucial determinant of the magnitude of the child penalty. To do so, we contrast women whose opportunity cost of time spent outside the labor market is very different. Assuming productivities in child-rearing activities do not vary too much, these differences in opportunity costs are well approximated by differences in post-childbirth potential hourly wages. However, post-childbirth potential hourly wages are not observed for mothers who choose to leave the workforce due to children. Additionally, these wages could themselves be affected by children-related labor supply decisions. This would be the case if time spent outside the labor market translates into a slower accumulation of labor market-specific human capital. As a result, we consider instead pre-childbirth hourly wages, averaged over several years. Because (potential) hourly wages are strongly correlated overtime when restricting to a single worker, this gives us a reasonable proxy of potential post-childbirth hourly wages that is affected by neither the sample selection nor the simultaneity bias. In the end, we estimate the heterogeneity of the consequences of childbirth along the distribution of pre-childbirth wages.

We consider the short-run (one-year) to long-run (ten-year) impacts on several labor outcomes: total labor earnings, hourly wages, and labor supply at both extensive and intensive margins of employment. Our empirical strategy embeds a nonlinear difference-in-difference framework within a nonparametric ranking of individuals along the hourly wage distribution; the latter aims precisely at depicting the heterogeneity in individual labor market trajectories along the wage distribution. Our treatment group consists of parents with n children, while our control group contains parents with exactly $n - 1$ children. Our difference-in-difference approach is nonlinear: it is set in a multiplicative form. This makes it possible to decompose, in an accounting sense, the causal effect of parenthood on labor earnings as the sum of adjustments on different margins of labor supply, plus the changes in the wage rate.

We apply this method to French administrative data, namely, the DADS panel, a comprehensive linked employer-employee dataset¹ that covers the period from 1995 to 2015 and contains information on individuals' labor earnings and paid hours. This panel is merged with the census data from the EDP, including longitudinal birth and marriage records at the individual level. Due to the richness of the dataset, we are able to consider the above-mentioned control and treatment groups at specific locations of the hourly wage distribution.

Our main results are as follows: (i) high-wage women experience much smaller labor earnings losses due to childbirth than do their lower-paid counterparts; (ii) they are much less likely to interrupt their careers or reduce their paid hours; (iii) importantly, the magnitudes of the latter two effects exhibit monotonic behavior along the hourly wage distribution; (iv) hourly wage losses appear rather homogeneous along the pre-childbirth hourly wage distribution; (v) the heterogeneity of child penalties with respect to the hourly wage distribution persists in the long run, up to 10 years after childbirth; and (vi) by contrast, men's children-related labor supply decisions are almost independent of their ranking in the wage distribution: they do not decrease their labor force attachment due to children.

We relate the observed monotonic patterns to the increasing opportunity cost of time spent outside the workforce. These results strongly suggest that women's *absolute* labor market productivity is a key determinant of mothers' children-related

¹Filling out the DADS form is a mandatory part of the process of paying payroll taxes.

labor supply decisions. Indeed, high-wage mothers who bear a high cost of career interruption are very unlikely to opt out of working or reduce their working hours. Conversely, those with much larger work disincentives because current hourly wages make child benefits that compensate for a reduction in labor supply worth considering² are more likely to leave the workforce or at least reduce their labor supply at the intensive margin.

These results are consistent with models in which families choose between the time mothers spend in the labor market and the time they spend on children when making their labor supply decision, but do not take this trade-off into account when it comes to fathers. To put it differently, families seem to consider market services as acceptable substitutes for childcare provided by mothers: their decisions seem to reflect a comparison between the income generated by the time spent by mothers in the labor market and the costs incurred by this time, particularly for childcare. On the other hand, families do not seem to perceive paternal contributions in the same way. This would explain why, in the existing literature, mothers of young children change their behavior in response to reforms that make child-related career breaks more or less costly (Piketty, 2005; Lequien, 2012; Joseph et al., 2013), while recent reforms that specifically target fathers have virtually no impact on their behaviors (Périvier and Verdugo, 2021).

We provide additional evidence to support the validity of our approach. First, the potential endogeneity of fertility decisions with respect to labor outcomes does not affect the identification of child penalties. Second, we show that the heterogeneity of these child penalties is not the mere consequence of mean reversion. Third, we address the issue of potential selection into treatment driving our estimates. Fourth, we rule out the possibility of the heterogeneity in pre-childbirth wages being a confounding factor by resorting to survey data: preferences towards childcare exhibit only limited heterogeneity along the wage distribution. Furthermore, our results are robust to controlling for human capital accumulation decisions that would plausibly be driven by preferences towards childcare, and may affect pre-childbirth hourly wages.

²additionally, e.g., at the minimum wage level career interruptions do not significantly affect future career prospects

Literature Childbirths tighten time constraints and shift women’s labor supply as well as labor market outcomes, which helps explain a substantial share of the gender pay gap as shown by, e.g., the seminal contributions on the “motherhood penalty” by [Waldfogel \(1995, 1997, 1998\)](#). Recent empirical evidence suggests that motherhood not only explains a large part of the gender gap in labor earnings but also accounts for a growing share of this gap in developed countries ([Kleven et al., 2019b](#)). More generally, childbirths have been shown to explain a significant share of the aggregate gender gap, though there is no consensus on the exact share or whether this contribution is increasing over time ([Bertrand et al., 2010](#); [Wilner, 2016](#); [Adda et al., 2017](#); [Juhn and McCue, 2017](#); [Kleven et al., 2019b](#)).

The most prominent contribution to child penalties likely stems from children-induced career interruptions and adjustments in labor supply, which in turn results in human capital depreciation ([Meurs et al., 2010](#); [Ejrnæs and Kunze, 2013](#); [Adda et al., 2017](#)). Other channels involve reduction in work effort ([Becker, 1985](#); [Hersch and Stratton, 1997](#)) and mothers having a strong preference for time flexibility ([Anderson et al., 2003](#); [Goldin, 2014](#)), which can generate compensating wage differentials or lead mothers to work in family-friendly firms that are likely to exert monopsony power ([Coudin et al., 2018](#)).

As to the causes of such decisions, two views can be contrasted. The first builds on the model of time allocation proposed by [Becker \(1981\)](#), based on the comparative advantage between the labor market and home production, i.e., on specialization. The second view, related to preferences and norms, refers to the identity model of [Akerlof and Kranton \(2000\)](#) and suggests that childbirth enhances the perception of oneself and her spouse as belonging to one gender or another, which distorts households’ time allocation decisions in the sense that is compatible with gender-specific prescriptions.

The empirical evidence seems to favor the second channel as the most compelling one. Indeed, cross-country comparisons ([Kleven et al., 2019a](#)), as well as comparisons of biological and adoptive families ([Kleven et al., 2021](#)), different- and same-sex couples ([Andresen and Nix, 2021](#)), or families belonging to different linguistic groups ([Steinhauer, 2018](#)), and lastly careful investigations of numerous family policy reforms ([Kleven et al., 2020](#)) all suggest that: (i) holding constant norms and preferences, differences in comparative advantage do not translate into differences

in child penalties; conversely, holding constant the comparative advantage, heterogeneity in exposure to different norms is strongly correlated with differences in child penalties. An exception to this trend would be [Angelov et al. \(2016\)](#) who find substantial heterogeneity in child penalties depending on parents' relative potential earnings.

This paper is also related to a few additional studies that have contrasted child penalties among individuals characterized by mothers' labor market opportunities. Firstly, following [Goldin \(2014\)](#), [Bütikofer et al. \(2018\)](#) compare child penalties across occupations among top earners. Secondly, a small body of sociological literature has been devoted to the distributional impact of the child penalty, following [Budig and Hodges \(2010\)](#). Due to methodological issues regarding the interpretation of quantile regression coefficients, it, however, remains difficult to identify the main lessons from this literature (see [Killewald and Bearak, 2014](#); [Budig and Hodges, 2014](#); [England et al., 2016](#)).

The above empirical strategies rely on an evaluation of the causal impact of parenthood on labor outcomes, which requires overcoming the issue of endogeneity of fertility decisions (see e.g. [Lundberg and Rose, 2000](#); [Miller, 2011](#)). However, [Kleven et al. \(2019b\)](#) emphasize that, empirically, correcting for potential endogeneity does not make too much of a difference. In this paper, we rely to some extent on this result to advocate for our difference-in-difference strategy, and develop additional tests that enable us to show that endogeneous fertility decisions likely do not affect our results.

Lastly, this paper is relevant to the analysis of heterogeneity of the gender pay gap along the wage distribution (e.g., [Albrecht et al., 2003](#); [Arulampalam et al., 2007](#); [Gobillon et al., 2015](#)). In particular, [Fortin et al. \(2017\)](#) points out that vertical segregation, i.e., women being underrepresented at the very top of the distribution, can account for a large share of the aggregate gender gap in earnings. Our results suggest that while child penalties may well contribute to this underrepresentation at the top, it is not the sole explanation: child penalties are, if anything, smaller at the top of the distribution. Yet vertical segregation may result from (even small) motherhood penalties: due to statistical discrimination, the generosity of parental leave systems may cause employers to place fewer women in top positions ([Datta Gupta et al., 2008](#); [Albrecht et al., 2015](#)).

The rest of the paper is organized as follows. The next section presents our data and the institutional setting. In section 1.3, we describe our empirical approach. Section 1.4 presents our results; section 1.5 discusses the validity of our identification strategy, and section 1.6 concludes.

1.2 Data and institutional background

1.2.1 Data

Our analysis is based on a large panel of French salaried employees, namely, the longitudinal version of the *Déclarations Annuelles de Données Sociales* (DADS). By law,³ French firms have to fill out the DADS form – an annual form that is the analogue of the W-2 form in the US – for every employee subject to payroll taxes. Starting from 1967, the panel covers individuals born in October of even-numbered years. As of 2002, the panel contains information on individuals born on January 2-5, April 1-4, July 1-4 and October 1-4 regardless of the parity of their year of birth; these (more or less) first four days of each quarter correspond to the birthdays of individuals for whom we obtain census records in addition to labor market characteristics. This panel is therefore a representative sample of the French salaried population at rate 4.4%. Because of the comprehensiveness of the panel with respect to individuals' careers, the data is of exceptional quality and has low measurement error in comparison with survey data, in addition to a large sample size and no top-coding.

The database contains detailed information about gross and net wages, days worked, paid hours,⁴ other job characteristics (the beginning, duration and end of a period of employment, seniority, and part-time employment), firm characteristics (industry, size, and region) and individual characteristics (age and gender). We are also able to recover the numbers of male and female employees at each firm by resorting to the cross-sectional version of the DADS to this end and using the linked employer-employee dataset (LEED) dimension. Our main variables of interest are (i) net real annual labor earnings defined as the sum of all salaried earnings over all employers, (ii) time worked, measured as the number of paid hours as well as

³The absence of DADS as well as incorrect or missing answers are punished with fines.

⁴This information has been available since 1995 only.

the number of days worked, and (iii) hourly wages defined as the ratio of annual earnings and time worked. In Appendix 1.A, we provide some further details on the measurement of earnings and time worked. The main point is that, with few exceptions, (i) maternity leave allowances paid by social security are not included in our measure of earnings; (ii) the duration of maternity leave in days corresponds to a positive number of days worked; (iii) the number of hours worked during the maternity leave is equal to 0, and (iv) the number of hours worked (resp., hourly wages) is overestimated (resp., underestimated) for workers that are not paid by the hour in years in which they take maternity leave.

Individuals are identified by their NIR, a 13-digit social security-like number that allows to merge the DADS panel with *Échantillon démographique permanent*. The latter is a longitudinal version of the census that includes births and marriage registers as of 1968. However, information on childbirth is missing before 2002 for individuals born in January, April or July. For this reason, we consider first individuals born on October 1-4. Additionally, some childbirth-related data is available in administrative birth registers for individuals born October 2-3; however, it was incomplete during the 1990s (for details, see Wilner, 2016): as a result, for these individuals we rely on the census rather than birth records.⁵ Finally, partial data on education is available in this dataset (see Charnoz et al., 2011) that indicates the highest degree obtained at the end of studies.

Our working sample is composed of salaried male and female employees in the private sector with the exclusion of agricultural workers and household employees. We restrict our analysis to individuals aged 20 to 60 living in metropolitan France between 1998 and 2015. This requires to restrict our attention to individuals born on even-numbered years, given that individuals born on odd-numbered years are not covered by the panel before 2002. We are therefore relying on a representative sample at rate 0.5%.⁶

The empirical analysis described in Section 1.3 requires selecting individuals with a strong attachment to the labor market. We specify that these individuals be

⁵Appendix 1.B explains how we recover such data, the quality of which is comparable with that of individuals born October 1 or 4 for whom birth records are available.

⁶On top of this longitudinal sample, we also rely on a comprehensive version of the DADS dataset that allows us to track all salaried employees from one year to the next to devise additional tests of our identifying assumption: see Subsection 1.5.1.

employed in the private sector for at least two years between $t-5$ and $t-2$ in addition to being present in $t-1$.⁷ To deal with individuals with very low labor participation, an individual is considered employed at t if her paid hours exceed 1/8 of the annual duration of work (1,820 hours as of year 2002), if her total employment duration exceeds 45 days per year and if her hourly wage exceeds 90% of the minimum wage. We also winsorize labor earnings at the quantile of order 0.99999 to avoid outliers. We exclude individuals for which one observation has the ratio of net labor earnings to gross labor earnings less than (resp., greater than) 1/100 of (resp., 100). Our working sample has approximately 1.4 million individuals-years of observations, corresponding to nearly 155,000 workers.

1.2.2 Summary statistics

Table 1.1 provides several statistics for the selection process. First, censoring of observations with low numbers of paid hours or low employment duration is illustrated. Second, the restriction to individuals for whom data are available for two years between $t-5$ and $t-2$ in addition to years $t-1$ and t is applied. As expected, both steps increase average hourly wages within a given gender, age group and industry. The selection is harsher for women than it is for men, as women are more likely to experience career interruptions. Censoring reduces the share of younger workers slightly, which is consistent with entry into the workforce through shorter and non-full-time employment spells; selection has the same effect for the same reason. Censoring reduces the share of workers in the service industry who are more likely to have short employment spells and to work part-time. Selection also reduces the share of service industry workers among men and the share of trade industry workers among women, as these individuals have less stable employment histories than those of their counterparts working in other industries.

Both within our base sample (after censoring) and within our selected sample, the gender gap in hourly wages is larger among older workers than among their younger counterparts.

Figure 1.1 displays the number of childbirths both in the raw EDP dataset and in

⁷The core results of this paper rely on years t from 1998 to 2015. As a result, because data are only available from 1995, the inclusion condition is slightly stronger for years 1998 and 1999. However, dropping these years and focusing only on years 2000 to 2015 does not change our estimates, as shown in Figures 1.E.9 and 1.E.10.

our final sample.⁸ Because we focus on childbirths that occur after individuals have experienced rather stable employment for several years in a row, and because our data only covers salaried employment in the private sector, numerous childbirths are not included in our final sample: we disregard about a half of women who experienced childbirth between 1998 and 2015. These proportion amount to roughly 60% for men during the same period.

1.2.3 Institutional background

Family-friendly policies in France have a long-lasting history (see [Rosental, 2010](#)) that dates back at least to pro-natalist concerns during the interwar period ([Huss, 1990](#)). These policies rely on (i) tax cuts, especially the *quotient familial* introduced in 1945, whereby the income tax rate depends on the number of children in a household, (ii) various child benefits, and (iii) some other welfare benefits, such as bonuses included in retirement pensions that depend on realized fertility, or housing allowances. In France, income is taxed jointly within households; this scheme is the source of strong incentives towards within-household specialization.

Maternity leaves were created in 1909; they were first unpaid, and subsequently became fully covered up to some threshold for all salaried workers by social insurance from 1970 onwards. Since 1980, the arrival of the first two children granted a woman a 16-week maternity leave consisting of 6 weeks before childbirth and 10 weeks after. Starting from the arrival of the third child, the total duration becomes 26 weeks (8+18), and maternity leave duration may increase to 46 weeks in the case of multiple births. Maternity leaves also have a minimum duration of 8 weeks, consisting of 2 weeks before childbirth and 6 weeks after. By contrast, paternity leaves have granted fathers an 11-day leave since 2002 only.

In addition to the above leaves, there are various parental allowances that were merged in 2004 into the *Prestation d'Accueil du Jeune Enfant* (PAJE).⁹ In fact,

⁸The raw EDP dataset itself is not perfectly representative of all childbirths that occur in France because it only provides information on fertility of individuals that have appeared at least once in labor market data, the sample of which has varied over time.

⁹It 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* (CMG)), and some child benefits granted when parents interrupt their careers or work part-time (previously *Complément Libre Choix d'Activité* (CLCA) and now *Prestation Partagée d'Éducation de l'enfant* (PreParE)).

these child benefits date back to 1985 and appeared with the creation of *Allocation Parentale d'Éducation* (APE) initially restricted to mothers of 3 or more children. APE was extended to mothers of 2 children in 1994, and was replaced by the CLCA in 2004, becoming effective with the first childbirth and providing a fixed not-means-tested amount for the maximum duration of 6 months. The CLCA was replaced in 2015 by PreParE that introduced incentives to split the leave between parents; it amounted to approximately €400 per month in the case of career interruption and to nearly €200 in the case of 80% part-time work. Several papers have shown that these benefits induce mothers to reduce their labor supply (Choné et al., 2004; Piketty, 2005; Lequien, 2012; Joseph et al., 2013).

In contrast, other policies favor participation in the labor force by decreasing the cost of childcare; an example of such a policy is CMG that is not means-tested, and entails payroll tax cuts or income tax credits.¹⁰ It is not straightforward to determine the exact scheme of financial incentives provided by such childcare subsidies because they depend on numerous dimensions (the type of childcare chosen among day nurseries, child-minder and nannies,¹¹ family structure and geographic location) but always depend on earnings in a way that makes mothers at the bottom of the wage distribution more likely to stop or reduce their activity (see, e.g., Givord and Marbot, 2015).

Considering labor supply, the current family insurance scheme therefore provides contradictory incentives: on the one hand, PreParE should reduce labor supply after childbirth; on the other hand, CMG should preserve it. Determining which effect dominates is an empirical task; yet the answer to that question depends crucially on the location in the wage distribution. Mothers at the top of the wage distribution will not be particularly responsive to PreParE since career interruption and part-time employment are more costly for them. In contrast, the combination of PreParE benefits (€200) with a reduction of childcare expenditures is worth considering for low-earnings mothers: e.g., at the minimum wage (slightly above €1,200 per month),

¹⁰A typical tax credit amounts to 50% of childcare expenditures up to some threshold that depends on the type of chosen daycare. The annual threshold is €2,300 for childcare providers or wet nurses, but it may be as high as €13,500 (€16,500 in the first year) for nannies employed at home.

¹¹This very choice itself depends on parents' earnings; affluent households are more likely to opt for nannies, while poor households more often choose child-minders or day nurseries, though there is variation in this respect.

a switch to 80% part-time work means a monthly cut of approximately €240, hence a net monetary loss of €40 only. Hence the current system including family allowances and childcare subsidies is more likely to make the “mommy track” all the more attractive to mothers located at the bottom of the wage distribution.

Other welfare benefits including bonuses included in pensions and housing allowances also depend on the number of children. Lastly, family-friendly policies may be available within firms; e.g., employers may provide childcare services to employees.¹²

1.3 Empirical analysis

Our main outcome of interest is total annual labor earnings of individual i during year t ; we denote such earnings by \tilde{y}_{it} . We decompose them into four components: d_{it} is a dummy variable for participation; \tilde{x}_{it} represents the employment duration in days, and is between 0 and 360;¹³ \tilde{h}_{it} denotes the average number of paid hours per day during year t , and lastly \tilde{w}_{it} is the average hourly wages of individual i during year t . Hence

$$\tilde{y}_{it} = d_{it}\tilde{x}_{it}\tilde{h}_{it}\tilde{w}_{it}. \quad (1.1)$$

1.3.1 Normalization

Providing estimates of the causal effect of childbirth by comparing parents and non-parents requires netting out other lifecycle effects as confounding factors; e.g., the number of childbirths an individual has experienced is a nondecreasing function of age. We choose to net out lifecycle and business cycle effects only; many other factors that determine labor outcomes could be adjusted in response to fertility decisions, and hence should be taken into account as part of child penalties instead of being controlled for. As a result, the first step of our empirical framework derived from that of [Guvenen et al. \(2021\)](#) consists of normalizing earnings and each of earnings’ components with respect to age, cohort and period. Let \tilde{z} denote either

¹²These firm-specific family policies can be subject to further tax reductions or credits, such as the *Crédit d’impôt famille* created in 2004.

¹³The number of days in a year is capped at 360 in DADS.

labor earnings or one of its components with the exception of the participation dummy. We start by regressing the logarithm of \tilde{z}_{it} on a set of cohort (year of birth), age and period dummies. We estimate the following pooled cross-sectional regression:

$$\log(\tilde{z}_{it}) = \sum_c \lambda_c^z \mathbf{1}_{cohort_i=c} + \sum_a \mu_a^z \mathbf{1}_{age_{it}=a} + \sum_T \nu_T^z \mathbf{1}_{t=T} + \epsilon_{it}^z \quad (1.2)$$

The identification of age-period-cohort (APC) models can be achieved at the cost of normalizations, which we detail in Appendix 1.C. In this paper, the choice of normalization is insignificant, given that we rely on the sum $\hat{\lambda} + \hat{\mu} + \hat{\nu}$ and never use these components separately.

Previous estimates enable us to define the normalized component z_{it} as

$$z_{it} = \frac{\tilde{z}_{it}}{\exp(\hat{\lambda}_{cohort_i}^z + \hat{\mu}_{age_{it}}^z + \hat{\nu}_t^z)} \quad (1.3)$$

An accounting decomposition similar to that of (1.1) is used for normalized earnings:

$$y_{it} = d_{it} x_{it} h_{it} w_{it} \quad (1.4)$$

1.3.2 Ranks in the hourly wage distribution

Our empirical strategy embeds a difference-in-difference setting within a framework that aims at modeling heterogeneity in the consequences of childbirth along the hourly wage distribution. To this end, we rely on comparisons both within groups of workers with similar hourly wages and across these groups. Hence our analysis relies on the definition of such groups based on a measure of recent hourly wages:

$$W_{i,t-1} = \frac{\sum_{\tau=t-5}^{t-1} d_{i\tau} \tilde{w}_{i\tau}}{\sum_{\tau=t-5}^{t-1} d_{i\tau} \exp(\hat{\lambda}_{cohort_i}^w + \hat{\mu}_{age_{i\tau}}^w + \hat{\nu}_\tau^w)} \quad (1.5)$$

We compute this measure for individuals who participate in year $t - 1$ and at least twice between years $t - 5$ and $t - 2$ (i.e., provided that $d_{i,t-1} \sum_{\tau=t-5}^{t-1} d_{i\tau} \geq 3$). Within each age \times year cell, we rank workers according to their recent wages $W_{i,t-1}$. We use this ranking to create 20 cells: P0-P5, P5-P10, ..., P90-P95 and P95-P100.

Hence we assume that workers within each age \times year \times recent wage cell are, if not identical, at least *ex ante* similar with respect to their hourly wage levels before year t . Ranks are not conditional on gender: within these cells, men and women have approximately the same recent wages. As a result, women are more (resp., less) numerous at the bottom (resp., top) of the distribution, which merely reflects the existence of a gender gap in hourly wages (see Table 1.1).

1.3.3 Difference-in-difference strategy

Our estimates of the consequences of childbirth are based on a difference-in-difference approach. The endogeneity of fertility decisions is often regarded as a key issue, but recent results suggest that it is not an empirical problem (Kleven et al., 2019b). We discuss the plausibility of the assumption that fertility decisions are exogenous, and devise additional tests of its validity in Section 1.5.

We define N treatments, where the n th treatment consists of experiencing the n th childbirth during year t . Our control group for the n th childbirth is composed of individuals of the same gender with $n - 1$ children and who never had the n th child. The main identifying assumption is that, absent the n th childbirth, the evolution of labor outcomes among parents of n children would have paralleled that of labor outcomes of parents that have exactly $n - 1$ children.

Due to the omission (“right-censoring”) of unknown but relevant data on fertility decisions taken after 2015, individuals belonging to the n th control group may experience the n th childbirth after 2015; we address this issue in Appendix 1.E. In practice, we restrict our attention to the first three childbirths that represent 96% of childbirths. Year $t - 1$ is regarded as the reference year; by construction, all individuals participate in the labor market during year $t - 1$.

Due to multiple treatments, the same individual may be considered several times in our estimation, though at different dates, either as a member of a treated group or a control group. Proper inference has to take this issue into account; we therefore cluster standard errors at the individual level (Bertrand et al., 2004).

This difference-in-difference approach is embedded in our ranking along the hourly wage distribution. Our control groups are therefore restricted to individuals with the same rank in the recent hourly wage distribution as our treated individuals. Moreover, the effect of childbirth is allowed to vary along that distribution of recent

wages.

The impact of the n th childbirth on earnings k years after childbirth for individuals of gender g at rank r in the recent wage distribution is given by

$$\beta_{g,r}^{y,n,k} = \underbrace{\log \left(\frac{\mathbb{E}[y_{i,t+k} | b_{it}^n = 1, r_{it} = r, g_i = g, t \in \mathcal{T}_k]}{\mathbb{E}[y_{i,t-1} | b_{it}^n = 1, r_{it} = r, g_i = g, t \in \mathcal{T}_k]} \right)}_{\text{Treated}} - \underbrace{\log \left(\frac{\mathbb{E}[y_{i,t+k} | c_{it}^n = 1, r_{it} = r, g_i = g, t \in \mathcal{T}_k]}{\mathbb{E}[y_{i,t-1} | c_{it}^n = 1, r_{it} = r, g_i = g, t \in \mathcal{T}_k]} \right)}_{\text{Control}} \quad (1.6)$$

where b_{it}^n is a dummy for experiencing the n th childbirth during year t , c_{it}^n is a dummy for belonging to the n th control group at time t , i.e., having $n - 1$ children at time t but never experiencing the n th childbirth according to the data, and \mathcal{T}_k is the set of time periods for which $t - 3$ to $t + k$ are observed in the data.¹⁴¹⁵ Notably, we do not match treatment and control groups according to age, period and cohort; this is made possible by our normalization of the data with respect to these dimensions that comes as a first stage in our empirical approach.

Considering the causal impact of childbirth $\beta_{g,r}^{y,n,k}$ being identified on a subset of time periods that depends on k , we assume that treatment effects are time-homogeneous, i.e., that having a k -year-old n th child bears the same consequences if the child was born in 1998 as it does if she was born in 2015. We assess the plausibility of this assumption, among others, in Section 1.5. Importantly, considering $k < -1$ allows us to verify that trends are parallel before childbirth.

The overall impact of childbirth on the gender gap in pay can be obtained directly as the difference between the impact on men's labor outcomes, and that on women's labor outcomes, both computed by the difference-in-difference method. It is thus written as (omitting indices for clarity):

$$\beta_{\text{gap}} = \beta_f - \beta_m. \quad (1.7)$$

¹⁴given the time-period that our dataset covers, this implies $\mathcal{T}_k = \llbracket 1998, 2015 - k \rrbracket$.

¹⁵Note that because these difference-in-difference estimates are directly obtained from the comparison of mean outcomes across groups and before and after childbirth, they are immune to concerns arisen by Goodman-Bacon (2021) and de Chaisemartin and D'Haultfœuille (2020) in the context of two-way fixed effects estimation of difference-in-difference models.

Decomposition (1.8) states that average normalized earnings growth can be represented as a sum of its four components, plus a selection term due to the fact that individuals who participate in the labor market in year $t+k$ may not have the exact same past earnings $y_{i,t-1}$ as those who do not participate:

$$\begin{aligned}
\underbrace{\log\left(\frac{\mathbb{E}[y_{i,t+k}]}{\mathbb{E}[y_{i,t-1}]}\right)}_{\text{Labor earnings changes}} &= \underbrace{\log(\mathbb{P}(d_{i,t+k} = 1))}_{\text{Participation}} \\
&+ \underbrace{\log\left(\frac{\mathbb{E}[y_{i,t-1}|d_{i,t+k} = 1]}{\mathbb{E}[y_{i,t-1}]}\right)}_{\text{Selection}} \\
&+ \underbrace{\log\left(\frac{\mathbb{E}[\mathbf{x}_{i,t+k}h_{i,t-1}w_{i,t-1}|d_{i,t+k} = 1]}{\mathbb{E}[\mathbf{x}_{i,t-1}h_{i,t-1}w_{i,t-1}|d_{i,t+k} = 1]}\right)}_{\text{Employment Duration Changes}} \\
&+ \underbrace{\log\left(\frac{\mathbb{E}[\mathbf{x}_{i,t+k}h_{i,t+k}w_{i,t-1}|d_{i,t+k} = 1]}{\mathbb{E}[\mathbf{x}_{i,t+k}h_{i,t-1}w_{i,t-1}|d_{i,t+k} = 1]}\right)}_{\text{Hours-per-day Changes}} \\
&+ \underbrace{\log\left(\frac{\mathbb{E}[\mathbf{x}_{i,t+k}h_{i,t+k}w_{i,t+k}|d_{i,t+k} = 1]}{\mathbb{E}[\mathbf{x}_{i,t+k}h_{i,t+k}w_{i,t-1}|d_{i,t+k} = 1]}\right)}_{\text{Hourly Wage Growth}}
\end{aligned} \tag{1.8}$$

This decomposition is made in an accounting sense.¹⁶ Specifically, a causal interpretation of this decomposition would require employment decisions to be mean independent of changes in the wage rate, which seems unlikely. In Appendix 1.D, we detail the computation of this decomposition, showing that it can be rewritten in terms of expected values of changes in labor outcomes, up to some reweighting. This decomposition of labor earnings growth allows us to consider separately each component of the impact of childbirth on earnings; we write it as $\beta^y = \beta^s + \beta^d + \beta^x + \beta^h + \beta^w$, where β^s stands for the selection term, and the four other terms correspond to each

¹⁶This decomposition is akin to the accounting decomposition of log-earnings changes as the sum of log-hourly wages and log-hours worked changes that is commonly used in labor economics (see e.g. Lachowska et al., 2020), while having the advantage of not conditioning on positive earnings. As a result, it allows to quantify, in an accounting sense, the contribution of the extensive margin of employment, which is particularly relevant in this particular setting.

component of labor earnings (for readability, we omit all other unnecessary indices).

1.4 Results

1.4.1 Heterogeneous consequences of childbirth

First, we assess the consequences of childbirth on labor outcomes of men and women by relying on the accounting framework. Our estimates of the impact of the first three childbirths on individuals' total labor earnings are shown in Figure 1.2 for women and in Figure 1.3 for men. We plot those estimates for $t+k \in \{t-3, \dots, t+5\}$ with the exception of $t-1$ since it is the reference year (our estimates are hence all equal to zero for that year).

Mothers experience large earnings losses after childbirth relative to women who earned similar hourly wages a few years before. In average, earnings losses due to the arrival of a first child amount to 30 log-points (26%) one to five years after her birth. All components contribute 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 are women belonging to our control groups. Nevertheless, in the short to medium run, labor supply decisions seem to be driving these large earnings losses. Moreover, the impact of childbirth on women's labor outcomes increases in magnitude with the rank of the child. This empirical evidence is consistent with previous findings in the literature.

More interestingly, children-related earnings losses display substantial heterogeneity: low-wage women experience far larger earnings losses than do high-wage women. At the very bottom of the distribution, women's earnings losses amount to 77 log-points (53%) the year of a child's arrival, and 33 log-points (28%) one year after childbirth, and remain at 52 log-points (40%) five years after the arrival of a child.¹⁷ In contrast, women ranked in the top 5% of the hourly wage distribution experience earnings losses of 22 log-points (20%), 9 log-points (8%) and less than 8 log-points (8%), respectively. The main result is that child penalties decrease along the wage distribution as pre-childbirth hourly wage increases.

¹⁷By definition and by law, year t includes a mixture of both maternity leave and employment periods, as discussed in Subsection 1.2.1.

The decomposition of annual earnings growth into each of its components helps clarify the channels that contribute the most to this pattern. Previous heterogeneity is primarily driven by labor supply decisions at the extensive margin: childbirth reduces by 17 log-points (resp., 60 and 80 log-points), i.e. 16% (resp. 45% and 55%) the probability that women are employed one year after the arrival of their first (resp., the second and the third) child at the bottom of the distribution but does not actually reduce this probability in the top 5% of the distribution. Once again, labor supply responses exhibit a striking monotonic behavior along the hourly wage distribution. Conversely, the pattern of motherhood wage penalties appears much more homogeneous one to five years after childbirth; the penalties amount to approximately 4 log-points (4%) for the first child, and even less for subsequent children.¹⁸

A nice feature of this approach is that it enables us to verify that trends of the treated and control groups before treatment are parallel. While observing parallel trends before treatment is not sufficient to assess the credibility of our identifying assumption,¹⁹ observing large differences in trends between treated and control groups before treatment would cast doubt as to the validity of our design. We observe small differences between groups' earnings in years $t-3$ and $t-2$ with respect to year $t-1$. The difference is slightly positive (resp., negative) when considering the arrival of the first (resp., the second) child: mothers had slightly slower (resp., faster) earnings growth than did non-mothers (resp., mothers of one child) prior to the first (resp., the second) childbirth. However, these differences are less than 10 log-points (9%), which is not much in comparison with earnings differences after childbirth (up to 130 log-points, i.e. 73%). More importantly, these differences vary little along the wage distribution, which is reassuring as far as the identification of heterogeneity of the impact of childbirth on women's labor outcomes is concerned.

When it comes to men, our estimates suggest that childbirths increase labor earnings slightly, especially through higher participation and hourly wages. The

¹⁸Hourly wage losses exhibit a U-shaped pattern along the distribution during the year of childbirth that may be driven by some problems in the measurement of hours during maternity leaves for workers that are not paid by the hour, who are more numerous in the upper part of the hourly wage distribution (see Subsection 1.2.1 and Appendix 1.A). Additionally, this U-shaped pattern may be due to the institutional setting: the maternity leave compensation scheme involves various thresholds and depends on its duration.

¹⁹This assumption deals with trends in potential outcomes (absent childbirth) after childbirth.

increase in participation is slightly more pronounced for fathers at the top of the wage distribution.

1.4.2 Long-run child penalties

To assess whether these heterogeneous child penalties persist in the long run, we extend our analysis to document the impact of children up to 10 years after childbirth. Figure 1.4 displays the corresponding estimates on labor earnings, participation, hours worked and hourly wages along the distribution. These estimates are based on a triple-difference approach: they correspond to the effect estimated for women minus that estimated for men. The figure makes it very clear that in the short to medium run, the gender gap in earnings widens much more at the bottom than at the top of the wage distribution. In the long run, it suggests that this large heterogeneity persists, as there is no convergence in child penalties 10 years after the arrival of the first child. This pattern is once again driven by differences in the impact on participation, though these results are less clear-cut than they are in the short to medium run, because standard errors grow large as time goes by. Furthermore, by focusing on the 1998-2015 time-period, we aggregate cohorts that were exposed to different policy environments, especially in terms of parental leave schemes.

1.5 Threats to identification

In this section, we address various issues that could affect the empirical validity of our identification strategy. Measurement error, including the right-censoring of fertility decisions, is considered in Appendix 1.E. The first issue has to do with the endogeneity of fertility decisions with respect to potential labor outcomes. The second concern is related to mean reversion while the third issue is potential selection into treatment. The last concern is the nonrandom assignment to pre-childbirth wage groups based on anticipated children-related labor supply decisions; this last source of endogeneity would not affect our estimates of the heterogeneity of child penalties *per se* but may rather affect their causal interpretation.

1.5.1 Endogeneity of fertility decisions

A first threat to identification stems from possible violations of the common trend assumption upon which our child penalties estimates are based. This assumption does not stand if individuals make their fertility decisions based on unobserved shocks common to both potential treated and untreated labor outcomes. Specifically, this would be the case if women expecting large earnings losses (due to dismissals for instance, or to cuts in the number of paid hours) to occur in the near future were more likely to have children. The parallel trend assumption would therefore not apply post-treatment, which would lead us to inflate the detrimental consequences of children.

In the absence of plausible exogeneous shocks to fertility decisions, there is no simple way of quantifying this potential source of bias. However, a recent empirical study by [Kleven et al. \(2019b\)](#) investigates this issue and observes that, for the third childbirth, child penalties estimated through simple event studies do not differ from those obtained by using a sex-mix instrument. Additionally, if high-wage women responded to expected future shocks to their labor outcomes the same way as low-wage women did, this source of bias would be constant along the distribution, and would not affect our claim that child penalties are larger at the bottom of the wage distribution than at the top.

In addition to these arguments, we provide direct evidence that plausible sources of negative shocks to labor outcomes do not trigger problematic fertility responses. First, we estimate how macro-level shocks in the labor market affect fertility decisions.

Within the population of eligible individuals, i.e., those with exactly $n - 1$ children at $t - 1$, we wonder how much b_{it}^n , the probability of birth of their n th child at time t , depends on the business cycle:

$$b_{it}^n = \eta^n \{\log(GDP_t) - \log(GDP_{t-1})\} + \kappa_{age_{i,t}}^n + \pi^n t + \xi_{it} \quad (1.9)$$

where all coefficients are indexed by rank in the recent wage distribution and gender. Coefficients η account for the sensitivity of fertility decisions to macro-level shocks. An endogeneity problem would arise if those coefficients were estimated to be significantly negative, especially at the bottom of the wage distribution. According to

Figure 1.5, with very few exceptions, this is not the case.

Second, we ask whether micro-level shocks generate such fertility responses. We build on [Huttunen and Kellokumpu \(2016\)](#), who show that job displacement triggers negative fertility responses. We rely on the linked employer-employee nature of our data to identify plausible mass layoff episodes. Namely, we assume that individual i is subject to a firm-level shock f_{it} at time t if more than 25% of individuals working for the main employer²⁰ of i at time $t - 1$, but who are not individual i herself, leave the firm at time t .²¹ Within each eligible subpopulation, these firm-level shocks indeed correlate with job losses. We estimate a linear model for the probability l_{it} of being jobless at time t ,

$$l_{it} = \rho^n f_{it} + \sigma_{age_{it},t}^n + v_{it}, \quad (1.10)$$

where all coefficients depend on gender and the rank in the wage distribution, omitting once again the index n . Figure 1.6 displays the estimates of coefficients ρ , and shows that it is plausible that exogenous firm-level shocks are felt at the individual level. We then estimate the probability of having the n th child at time t ,

$$b_{it}^n = \phi^n f_{it} + \psi_{age_{it},t}^n + \omega_{it} \quad (1.11)$$

Figure 1.7 displays the corresponding estimates of ϕ . In most cases, we cannot reject the null hypothesis that these coefficients are equal to 0, which suggests that firm-level employment shocks do not trigger positive fertility responses that would render our estimates of child penalties meaningless.

1.5.2 Mean reversion

We address next the issue of mean reversion. Our approach that aims at comparing child penalties across the hourly wages distribution requires splitting the sample according to pre-event hourly wages. A potential shortcoming of this approach is that, by doing so, we may end up introducing mean reversion in the estimation: if initial

²⁰The main employer of an individual is defined as the firm that pays that individual the largest earnings during a given year.

²¹These firm-level shocks are identified from a comprehensive version of the DADS data that allows to track all salaried employees, and not only those included in the DADS panel, from year $t - 1$ to year t .

labor outcomes incorporate both a permanent and a transitory component, then women with high initial earnings would be mechanically more likely to experience relative earnings fall in subsequent years. In other words, the assumption that we compare women with different opportunity costs fails if we condition on short-term shocks.

To check that this does not drive our estimates, we first display changes between the last year before the (counterfactual) childbirth in both treatment and control groups, and the (counterfactual) childbirth year. Figure 1.8 display our results. Firstly, it makes it clear that heterogeneity in the child penalties stems from the heterogeneity in the earnings changes in the treated group, not in the control group as would be the case if mean reversion were at play. Secondly, the heterogeneity in the earnings changes of the control group does not suggest that we are conditioning on short-term shocks that would entail mean reversion: individuals at the bottom of the distribution are more likely to experience earnings losses, which is consistent with them being more exposed to unemployment risk (Güvenen et al., 2021).

To further insure against mean reversion driving our estimates, we replicate our approach while conditioning on hourly wages as observed 3 to 7 years before childbirth, as opposed to 1 to 5 years before childbirth (Figures 1.F.1 and 1.F.2). The results remain completely similar to our baseline estimates.

1.5.3 Selection into treatment

The next worrying issue is selection into treatment. In the presence of heterogeneous treatment effects, difference-in-difference estimates identify the average treatment effect on the treated (ATT). As a result, the estimates in the presence of subpopulations may differ due to two distinct channels: (i) in the labor market, low-wage women experience more detrimental consequences of child arrivals, and (ii) high-wage women – as defined by their *potential hourly wages absent children* – suffer from the same detrimental consequences of fertility but among them, those who face the largest career costs choose not to have children.

We can assess the plausibility of the second channel by computing the probability of the n th child arrival during year t among those eligible, i.e., those who already have $n - 1$ children in year $t - 1$, along the entire recent wage distribution. Figure 1.9 shows that among women, this probability does not vary much along the wage

distribution, and that high-wage women are, if anything, in fact more likely to have children than are their low-wage counterparts.

Lastly, our estimates of the consequences of the n th child aggregates the effect of this childbirth with that of possible subsequent childbirths. Mothers at the top of the distribution might therefore face smaller child penalties because, conditional on having one child, they would be less likely to have a second child. However, Figure 1.9 suggests that it is actually the reverse. To emphasize this point, Appendix 1.G displays the probability of experiencing a second childbirth a few years later for individuals that just had their first child. Parents at the top of the distribution are more likely to have a second child, hence the smaller child penalty is not driven by them limiting their subsequent fertility.

1.5.4 Endogeneity of pre-childbirth wages

After ruling out previous concerns, some doubts could still remain as to the random assignment of individuals to wage groups.²² Pre-childbirth wages may reflect individuals' decisions that likely depend on their preferences or on the gender norms they are exposed to.

We assess the extent to which that channel is likely to explain our results (i) by resorting to survey data in which individuals report their preferences towards childcare, in which we observe that the correlation between preferences and hourly wages is quite limited; (ii) by reweighting the data so that the implied composition in terms of past human capital investments and work-family preferences is the same all along the hourly wage distribution.

We resort to the French Labor Force Survey (LFS) which was complemented in 2010 with a module devoted to work-family balance. Individuals with children aged 3 or less were asked what was the ideal childcare solution to them for children as old as their youngest child. In Figure 1.10, we display the share of individuals declaring children should be taken care of by their parents along the hourly wage distribution. First, the vast majority of parents do not view parents-provided childcare as the best childcare solution for children aged the same age as their youngest child:

²²The ideal design for documenting the impact of absolute labor market productivities, for instance, would be to assign individuals exogenously to various productivities, i.e., to varying wage levels, and to contrast child penalties across those groups.

over 2/3 of parents favor external childcare, provided either in formal or informal settings (e.g. by grandparents). Second, there is limited heterogeneity along the wage distribution: the share of women (men) considering parents as the ideal childcare solution varies between 14% and 32% (23% and 48%). More importantly, the pattern is non-monotonic, which contrasts with our finding as to the profile of labor supply responses to childbirth along the hourly wage distribution, which suggests preferences may not be the main driver of these labor supply decisions. Additionally, the pattern for men does not match the pattern for women.²³

We then show that additional sources of heterogeneity, due to past human capital decisions which could affect pre-childbirth hourly wages and stem from childcare-related preferences, do not drive our results. To this end, we consider (i) education, measured by the highest degree obtained at the end of studies, as an 8-level variable; (ii) recent labor supply at all margins between year $t - 5$ and $t - 1$; (iii) the share of females working part-time for the main employer of each individual at time $t - 1$; and (iv) the age at (counterfactual) childbirth. We rely on these variables to reweight the data so that within each treatment/control group, the composition does not vary across the recent wage distribution. In this setting, the weight of the observation that corresponds to individual i at time t , who belongs to the treatment (control) group g , and is ranked r in the hourly wage distribution writes:

$$p_{it} = \frac{\mathbb{P}(R = r | G = g)}{\mathbb{P}(R = r | G = g, X = x_{it})} \quad (1.12)$$

where x_{it} corresponds to the observed variables upon which our reweighting procedure is based. Specifically, we take as $\mathbb{P}(R = r | G = g, X = x_{it})$ the predicted probability of belonging to rank r in the distribution based on an ordered logit.

We then replicate our analysis on the reweighted data. Our estimates represent heterogeneity in child penalties in a counterfactual population in which rank in the wage distribution would not be related to education, past labor supply, firm choice and age at childbirth. Figure 1.11 displays our results. While there might be slightly less heterogeneity along the wage distribution than in our baseline estimates, the

²³This descriptive evidence is affected by selection bias, given that hourly wages in the LFS can only be computed for individuals who are salaried employees; Appendix 1.H provides further evidence that addresses this issue. These results cannot reject the null hypothesis that parents with the most conservative views regarding childcare do not differ from the others in terms of their (potential) hourly wages.

patterns are still extremely similar.

1.6 Conclusion

This paper investigates whether mothers with different labor market opportunities make different children-related labor supply decisions, which in the end translate into earnings losses. To do so, we contrast the causal effect of children, identified thanks to a difference-in-difference approach, along the pre-childbirth wage distribution. We show that while, regardless of their wages, children have a large and negative impact on mothers' labor earnings, the magnitude of this impact is much larger for those with low potential hourly wages than it is for those with high potential hourly wages. The reason for this is that the former are much more likely than the latter to retrieve for the workforce, and to decrease their hours worked in the labor market. By contrast, fathers are very unlikely to change their hours worked upon the arrival of children, regardless of the wage rate.

Assuming that productivities in child-rearing activities are reasonably homogeneous, differences in potential hourly wages translate directly into differences in the opportunity cost of time spent outside the labor market. As a result, we interpret our findings as reflecting the fact the mothers' opportunity cost of time spent outside the labor market is a key determinant of children-related labor supply decisions, and as such, of the child penalty. This opportunity cost sums up the trade-off between the income generated by the time mothers spend on the market and the costs incurred by this time, i.e. mothers' foregone contribution to child-rearing. High-wage mothers being much less likely to decrease their labor supply than their low-wage counterparts thus suggests that the former can compensate the latter. In other words, families may be willing to resort to market solutions that substitute for maternal care, provided that the cost remains sufficiently low. By contrast, fathers' labor supply decisions seem almost independent of this cost. This would indicate that in their time allocation problem, families do not view fathers' contribution to child-rearing as a possible substitute for mothers' contribution. Overall, this interpretation helps rationalize why, in the existing literature, mothers of young children respond strongly to reforms that make child-related career breaks more or less costly (Piketty, 2005; Lequien, 2012; Joseph et al., 2013), while recent reforms

that specifically target fathers have close to no impact on their behaviors ([Pérvier and Verdugo, 2021](#)).

References

- Adda, J., Dustmann, C., and Stevens, K. (2017). The career costs of children. *Journal of Political Economy*, 125(2):293–337.
- Akerlof, G. A. and Kranton, R. E. (2000). Economics and Identity. *The Quarterly Journal of Economics*, 115(3):715–753.
- Albrecht, J., Björklund, A., and Vroman, S. (2003). Is there a glass ceiling in sweden? *Journal of Labor Economics*, 41(1):89–114.
- Albrecht, J., Thoursie, P. S., and Vroman, S. (2015). Parental Leave and the Glass Ceiling in Sweden. In *Gender Convergence in the Labor Market*, volume 41 of *Research in Labor Economics*, pages 89–114. Emerald Publishing Ltd.
- Anderson, D. J., Binder, M., and Krause, K. (2003). The motherhood wage penalty revisited: Experience, heterogeneity, work effort, and work-schedule flexibility. *Industrial and Labor Relations Review*, 56(2):273–294.
- Andresen, M. E. and Nix, E. (2021). What Causes the Child Penalty? Evidence from Adopting and Same-Sex Couples. mimeo.
- Angelov, N., Johansson, P., and Lindahl, E. (2016). Parenthood and the Gender Gap in Pay. *Journal of Labor Economics*, 34(3):545–579.
- Arulampalam, W., Booth, A. L., and Bryan, M. L. (2007). Is there a glass ceiling over europe? exploring the gender pay gap across the wage distribution. *Industrial and Labor Relations Review*, 60(2):163–186.
- Becker, G. (1981). *A Treatise on the Family*. Harvard University Press, Cambridge.
- Becker, G. S. (1985). Human Capital, Effort, and the Sexual Division of Labor. *Journal of Labor Economics*, 3(1):33–58.
- Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How Much Should We Trust Differences-In-Differences Estimates? *The Quarterly Journal of Economics*, 119(1):249–275.
- Bertrand, M., Goldin, C., and Katz, L. F. (2010). Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, 2(3):228–55.
- Budig, M. J. and Hodges, M. J. (2010). Differences in disadvantage: Variation in the motherhood penalty across white women’s earnings distribution. *American Sociological Review*, 75(5):705–728.
- Budig, M. J. and Hodges, M. J. (2014). Statistical models and empirical evidence for differences in the motherhood penalty across the earnings distribution. *American Sociological Review*, 79(2):358–364.
- Bütikofer, A., Jensen, S., and Salvanes, K. G. (2018). The role of parenthood on the gender gap among top earners. *European Economic Review*, 109:103–123.
- Charnoz, P., Coudin, E., and Gaini, M. (2011). Changes in the french wage distribution 1976-2004: Inequalities within and between education and experience groups. Working paper INSEE.
- Choné, P., Le Blanc, D., and Robert-Bobée, I. (2004). Offre de travail féminine et garde des jeunes enfants. *Économie et Prévision*, 162(1):23–50.

- Coudin, E., Maillard, S., and Tô, M. (2018). Family, firms and the gender wage gap in france. IFS Working Papers W18/01, Institute for Fiscal Studies.
- Datta Gupta, N., Smith, N., and Verner, M. (2008). The impact of nordic countries' family friendly policies on employment, wages, and children. *Review of Economics of the Household*, 6(1):65–89.
- de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–96.
- Deaton, A. S. (1997). Econometric issues for survey data. In Bank, T. W., editor, *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*, pages 63–132. The Johns Hopkins University Press.
- Deaton, A. S. and Paxson, C. (1994). Saving, Growth, and Aging in Taiwan. In *Studies in the Economics of Aging*, NBER Chapters, pages 331–362. National Bureau of Economic Research, Inc.
- Ejrnæs, M. and Kunze, A. (2013). Work and wage dynamics around childbirth. *The Scandinavian Journal of Economics*, 115(3):856–877.
- England, P., Bearak, J., Budig, M. J., and Hodges, M. J. (2016). Do highly paid, highly skilled women experience the largest motherhood penalty? *American Sociological Review*, 81(6):1161–1189.
- Fortin, N. M., Bell, B., and Böhm, M. (2017). Top earnings inequality and the gender pay gap: Canada, sweden, and the united kingdom. *Labour Economics*, 47(Supplement C):107 – 123.
- Givord, P. and Marbot, C. (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., Meurs, D., and Roux, S. (2015). Estimating gender differences in access to jobs. *Journal of Labor Economics*, 33(2):317–363.
- Goldin, C. (2014). A Grand Gender Convergence: Its Last Chapter. *The American Economic Review*, 104(4):1091–1119.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*.
- Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021). What do data on millions of u.s. workers reveal about lifecycle earnings dynamics? *Econometrica*, 89(5):2303–2339.
- Hersch, J. and Stratton, L. S. (1997). Housework, fixed effects, and wages of married workers. *The Journal of Human Resources*, 32(2):285–307.
- Huss, M.-M. (1990). Pronatalism in the inter-war period in france. *Journal of Contemporary History*, 25(1):39–68.
- Huttunen, K. and Kellokumpu, J. (2016). The effect of job displacement on couples fertility decisions. *Journal of Labor Economics*, 34(2):403–442.
- Joseph, O., Pailhé, A., Recotillet, I., and Solaz, A. (2013). The economic impact of taking short parental leave: Evaluation of a french reform. *Labour Economics*, 25:63 – 75.

- Juhn, C. and McCue, K. (2017). Specialization then and now: Marriage, children, and the gender earnings gap across cohorts. *Journal of Economic Perspectives*, 31(1):183–204.
- Killewald, A. and Bearak, J. (2014). Is the motherhood penalty larger for low-wage women? a comment on quantile regression. *American Sociological Review*, 79(2):350–357.
- Kleven, H., Landais, C., Posch, J., Steinhauer, A., and Zweimüller, J. (2019a). Child penalties across countries: Evidence and explanations. *AEA Papers and Proceedings*, 109:122–26.
- Kleven, H., Landais, C., Posch, J., Steinhauer, A., and Zweimüller, J. (2020). Do family policies reduce gender inequality? evidence from 60 years of policy experimentation. Working Paper 28082, National Bureau of Economic Research.
- Kleven, H., Landais, C., and Søgaaard, J. E. (2019b). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.
- Kleven, H., Landais, C., and Søgaaard, J. E. (2021). Does biology drive child penalties? evidence from biological and adoptive families. *American Economic Review: Insights*, 3(2):183–98.
- Lachowska, M., Mas, A., and Woodbury, S. A. (2020). Sources of displaced workers' long-term earnings losses. *American Economic Review*, 110(10):3231–66.
- Lequien, L. (2012). The impact of parental leave duration on later wages. *Annals of Economics and Statistics*, (107/108):267–285.
- Lundberg, S. and Rose, E. (2000). Parenthood and the earnings of married men and women. *Labour Economics*, 7(6):689–710.
- Mason, K. O., Mason, W. M., Winsborough, H. H., and Poole, W. K. (1973). Some methodological issues in cohort analysis of archival data. *American Sociological Review*, 38(2):242–258.
- Meurs, D., Pailhé, A., and Ponthieux, S. (2010). Child-related career interruptions and the gender wage gap in france. *Annals of Economics and Statistics*, pages 15–46.
- Miller, A. R. (2011). The effects of motherhood timing on career path. *Journal of Population Economics*, 24(3):1071–1100.
- 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*, (156):79–109.
- Pora, P. and Wilner, L. (2020). A decomposition of labor earnings growth: Recovering gaussianity? *Labour Economics*, 63:101807.
- Périer, H. and Verdugo, G. (2021). Can parental leave be shared? OFCE Working Paper 6, OFCE.
- Rosental, P.-A. (2010). Politique familiale et natalité en france : un siècle de mutations d'une question sociétale. *Santé, Société et Solidarité*, 9(2):17–25.
- Steinhauer, A. (2018). Working Moms, Childlessness, and Female Identity. Sciences Po publications 79, Sciences Po.

- Waldfogel, J. (1995). The price of motherhood: Family status and women's pay in young british cohort. *Oxford Economic Papers*, 47(4):584–610.
- Waldfogel, J. (1997). The effect of children on women's wages. *American Sociological Review*, 62(2):209–217.
- Waldfogel, J. (1998). Understanding the "family gap" in pay for women with children. *The Journal of Economic Perspectives*, 12(1):137–156.
- Wilner, L. (2016). Worker-firm matching and the parenthood pay gap: Evidence from linked employer-employee data. *Journal of Population Economics*, 29(4):991–1023.

Figures

Figure 1.1 – Consequences of sample selection with respect to childbirths

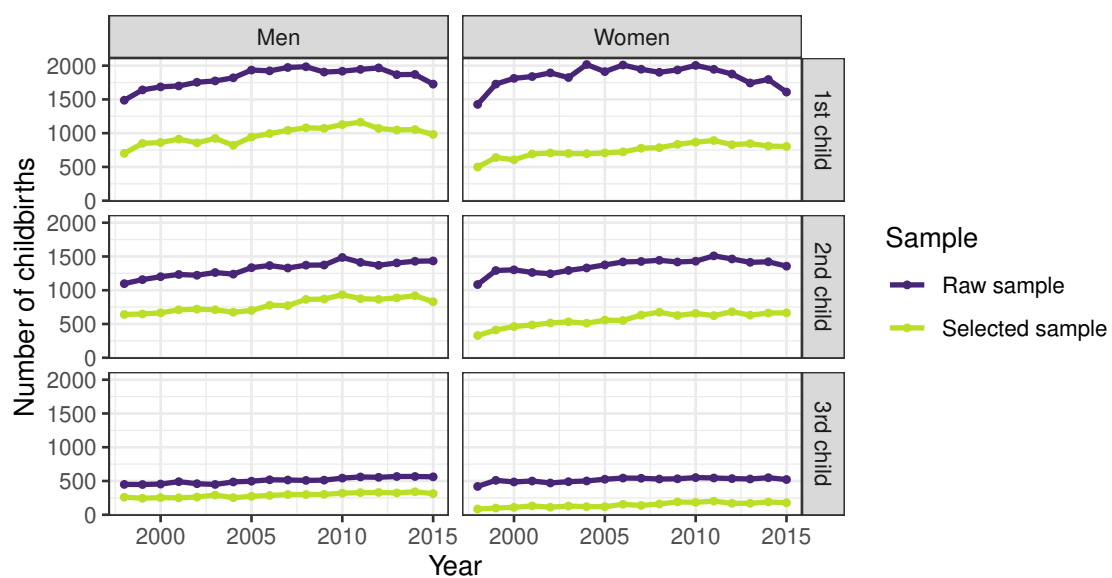

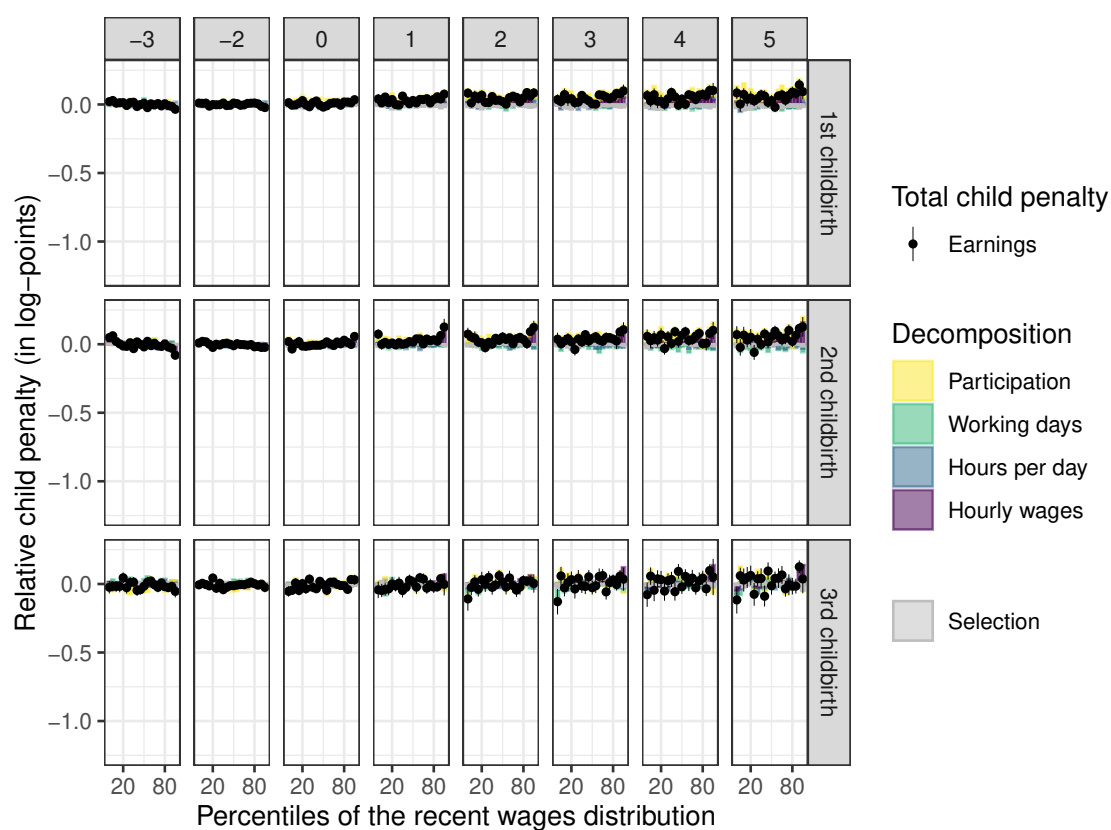


Figure 1.2 – Consequences of childbirth for women’s labor outcomes




Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see Equation (1.6)) for various values of time-to-childbirth expressed in years. Bootstrapped standard errors using 100 replications are clustered at the individual level.

Figure 1.3 – Consequences of childbirth for men’s labor outcomes




Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see Equation (1.6)) for various values of time-to-childbirth expressed in years. Bootstrapped standard errors using 100 replications are clustered at the individual level.

Figure 1.4 – Impact of first childbirth on the gender gap in earnings and labor outcomes




Estimates of the impact of first childbirth on the gender gap in pay, obtained by the difference-in-difference-in-difference method (see Equation (1.7)). Bootstrapped standard errors using 100 replications are clustered at the individual level.

Figure 1.5 – Probability of having children (sensitivity to the business cycle)



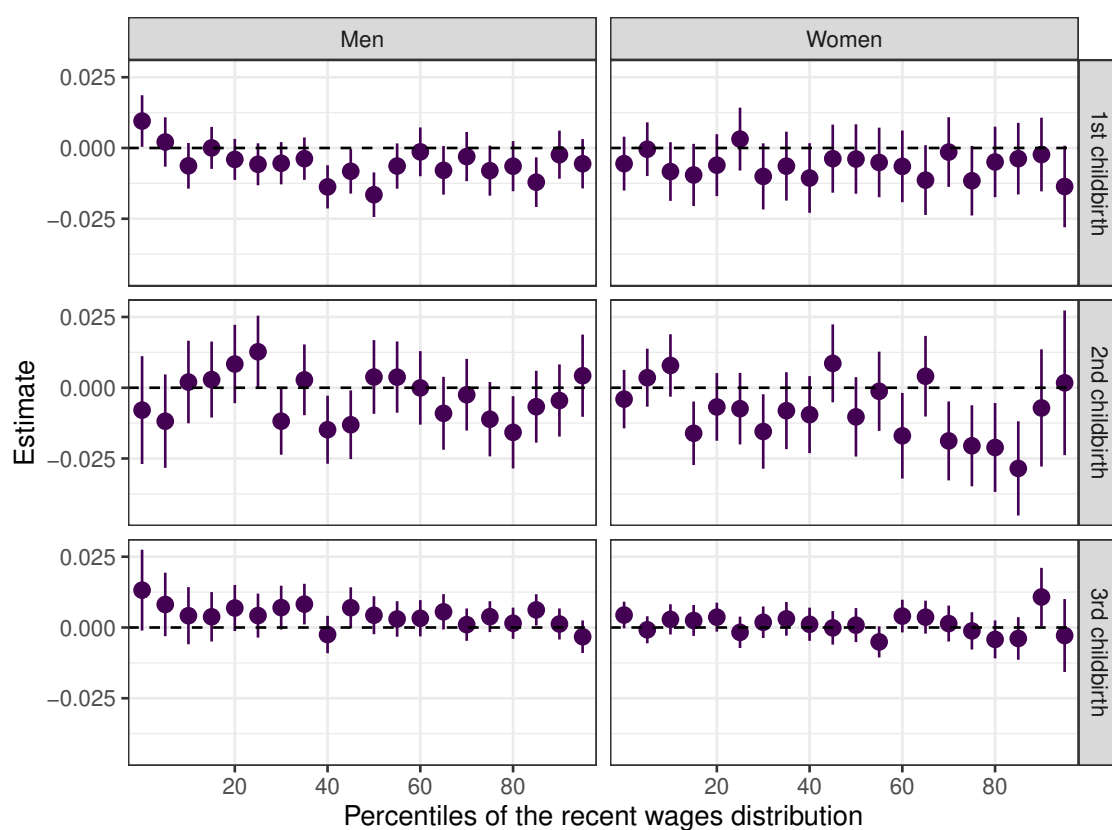
Estimates of coefficients related to log-GDP growth between times $t-1$ and t in a linear probability model with rank in the recent wage distribution \times age fixed effects (1.9). The outcome is a dummy variable for having the n th childbirth at time t . Standard errors are clustered at the individual level. The sample includes individuals up to age 60 at time t .

Figure 1.6 – Probability of losing one’s job (sensitivity to firm-level shocks)



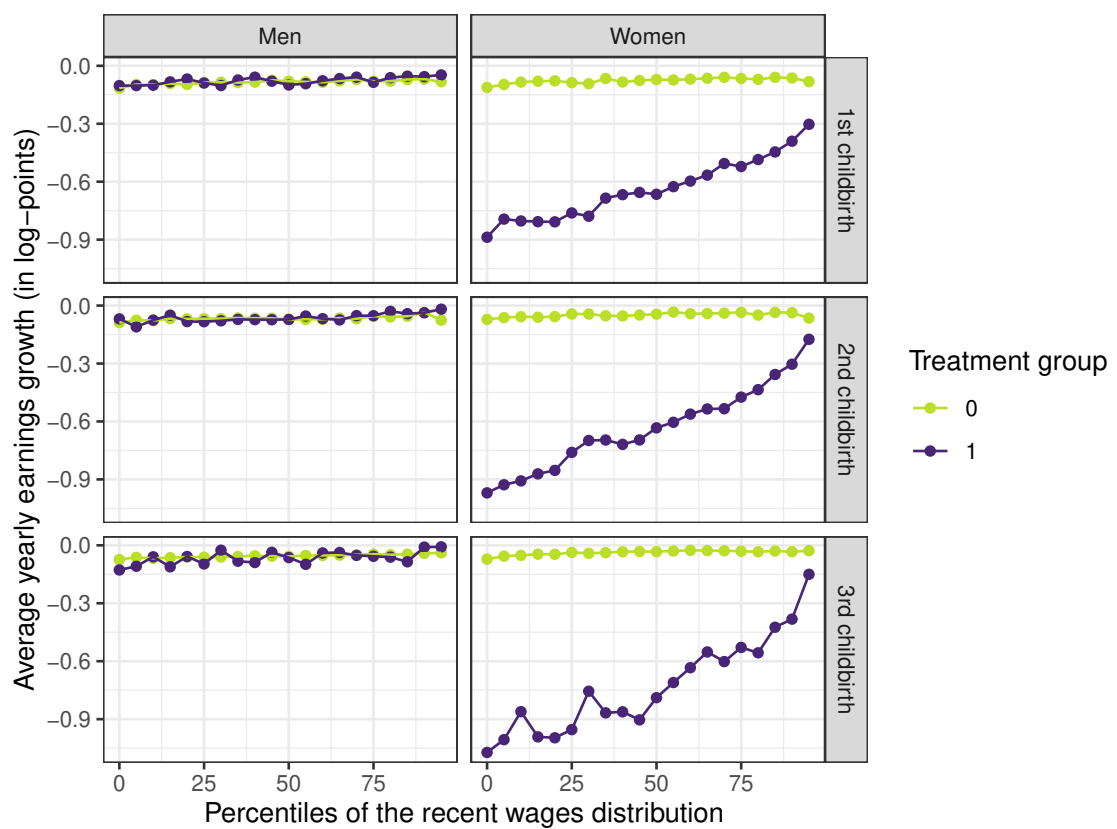
Estimates of coefficients related to firm-level shocks in a linear probability model with rank in the recent wage distribution \times age \times year fixed effects (1.10). The outcome is a dummy variable for being jobless at time t . Standard errors are clustered at the individual level. The sample includes individuals up to age 60 at time t .

Figure 1.7 – Probability of having children (sensitivity to firm-level shocks)



Estimates of coefficients related to firm-level shocks in a linear probability model with rank in the recent wage distribution \times age \times year fixed effects (1.11). The outcome is a dummy variable for having the n th childbirth at time t . Standard errors are clustered at the individual level. The sample includes individuals up to age 60 at time t .

Figure 1.8 – Underlying earnings changes



Log-average earnings changes between $t - 1$ and t by gender and treatment status, used in Figures 1.2 and 1.3.

Figure 1.9 – Probability of having children (by subperiod)

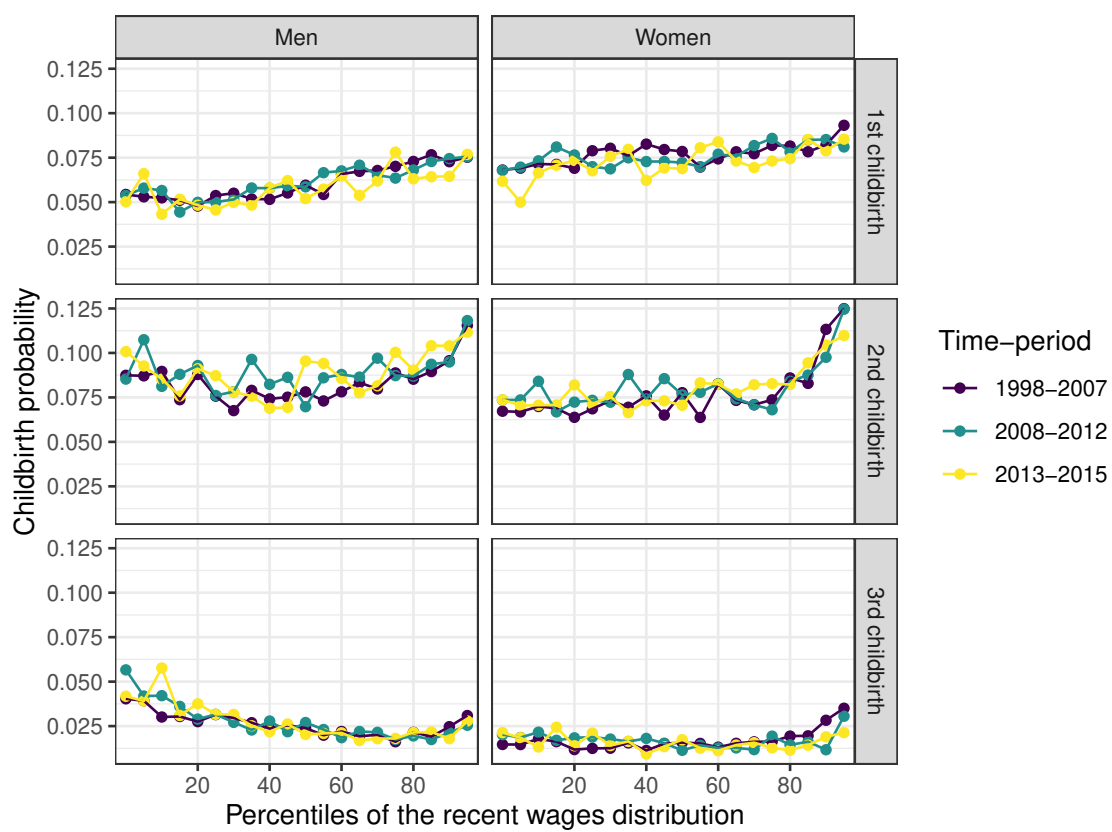




Figure 1.10 – Childcare preferences and hourly wages

Share of individuals who declare parents-provided childcare is the best childcare solution for children aged the same age as their youngest child. Ranks in the hourly wages distribution are conditional on the age of the youngest child. The sample includes individuals with children aged 3 or less.

Figure 1.11 – Counterfactual child penalties based on the reweighted data



Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see Equation (1.6)) for various values of time-to-childbirth expressed in years, based on the reweighted dataset in which variation along the wage distribution in education, past labor supply and sorting is cancelled out. Bootstrapped standard errors using 100 replications are clustered at the individual level.

Tables

Table 1.1 – Sample selection

	Base sample				Censoring				Final sample			
	Women		Men		Women		Men		Women		Men	
# Obs.	907 202		1 166 815		778 320		1 062 183		581 921		832 477	
# Ind.	97 388		110 818		88 015		104 322		68 235		87 148	
	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	7	9,2	6,2	9,5	6,2	9	5,5	9,5	4,4	9,2	4	9,8
25-29	16,9	10,3	16,1	11,3	16,5	10,3	15,8	11,1	15,4	10,5	14,6	11,4
30-34	15,3	11,6	15,7	13,2	15,2	11,4	15,8	12,9	15,1	11,8	15,8	13,2
35-39	14,7	12,2	15,2	15,2	14,9	12	15,4	14,4	15	12,5	15,7	14,7
40-44	14	12,5	14	16	14,3	12,3	14,3	15,5	14,7	12,8	14,7	15,9
45-49	13,4	12,7	13,2	16,5	13,7	12,5	13,4	16,3	14,4	12,9	13,9	16,6
50-54	11	12,8	11,4	17,5	11,4	12,8	11,6	17,3	12,4	13,1	12,3	17,5
55-59	7,6	13,7	8,2	18,7	7,8	13,2	8,3	18,4	8,7	13,4	8,9	18,8
Industry												
Construction	1,7	12,1	11,5	12,9	1,9	12,3	12	12,6	2	12,9	12,3	13,3
Manufacturing	13,7	12,3	25,3	15,1	14,8	12,1	26,8	14,8	15,7	12,6	28,4	15,2
Services	64,8	11,8	47,6	14,5	62,9	11,8	45,1	14,6	61,8	12,5	43,2	15,6
Trade	19,8	10,5	15,6	13,9	20,4	10,4	16	13,3	20,4	10,9	16,2	14

Base sample includes all individuals aged 20 to 60 that have positive employment in the private sector at time t . Censoring excludes individuals that that work less than 45 days a year, less than 1/8 of the legal duration a week, or paid less than 90% of the minimum hourly wage. Final sample includes only individual that are over this threshold at time t , $t - 1$ and at least twice between $t - 5$ and $t - 5$. Figures for the final sample are computed at time $t - 1$.

1.A Earnings and working time measures

1.A.1 Earnings

Our measure of labor earnings relies on net annual earnings. This measure aggregates all wages paid to an individual, including performance pay and bonuses, paid vacations, in-kind benefits, the share of severance payments that exceeds the legal minimum, and early retirement benefits (to the extent that these benefits exceed an amount approximately equal to the minimum wage) but excludes stock-options. Social security contributions, public pension schemes, unemployment benefits and other contributions including two flat-rate taxes on earned income (CSG and CRDS) are subtracted to this amount to compute our measure of net annual earnings. In that sense, we measure earnings before income taxes but after some transfers.

Maternity leave allowances are paid by the Social Security administration, and as such are not part of our measure of earnings. They may, however, be paid through the employer (*subrogation*): in this setting, the employer pays the employee the equivalent of maternity leave allowances during her maternity leave, and is later reimbursed by the Social Security administration. The employer subsequently subtracts the maternity leave allowances that the employer advanced from the measure of earnings. Because the reimbursement occurs after the maternity leave itself, the decline in earnings may occur a few weeks after the maternity leave. Because we consider annual earnings, this problem is restricted to childbirths that occur at the end of the calendar year. Our results are, however, very robust to considering only childbirths that occur in the 2nd quarter of the year that are immune to this issue (see Appendix 1.E).

Lastly, in some firms the employer may be bound by collective agreement to complement earnings during maternity or sick leaves in addition to Social Security-provided allowances. This complement is part of labor earnings as measured by the DADS.

1.A.2 Days

In the DADS dataset, days worked refer to the duration during which an employee is part of the workforce of a firm within a given year. As a result, maternity and sick leaves, or paid vacations are part of this measure of days, whereas a period of unemployment between two distinct employment spells is not. Additionally, this

measure of days is capped at 360.

1.A.3 Hours

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

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

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

As to workers whose compensation does not depend on the time worked, but who do not belong to one of the above-mentioned occupations, i.e., typically managers ("forfait-jour"), employers provide the number of days only. A number of hours is ascribed to these observations based on the legal duration of work for full-time workers, the number of work days, and the implied hourly wages.

Because during a maternity leave, an employee is not paid by her employer for any hours worked but is instead paid by the Social Security Administration (and possibly receives a complementary payment from her employer), hours worked during a maternity leave are equal to 0. Workers who are not paid by the hour are an exception to this rule because their hours are imputed based on their days worked, which do not vary during maternity leaves. As a result, the DADS dataset overestimates hours worked – and underestimates hourly wages – for such workers during years when they give birth to children. In general, these are qualified workers that belong to the upper part of the hourly wage distribution, so the decomposition

of earnings penalties into hours and hourly wages may be biased at the top of the distribution for the specific year workers take maternity leaves.

1.B Childbirth imputation

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

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

- the years of birth of the 1st to the 12th children appearing in birth records as of 1967;
- the years of birth of the 1st to the 12th children as declared in the 1990 census;
- the years of birth of the 1st to the 12th children as declared in the 1999 census.

Information from birth records has been available since 1967 only, which results in left-censoring. However, because we are mostly interested in individuals giving birth between 1998 and 2015, we do not try to deal with this issue. Our goal is to fill the gap in administrative records between 1982 and 1997 for half of the sampled individuals, which increases our sample size substantially.

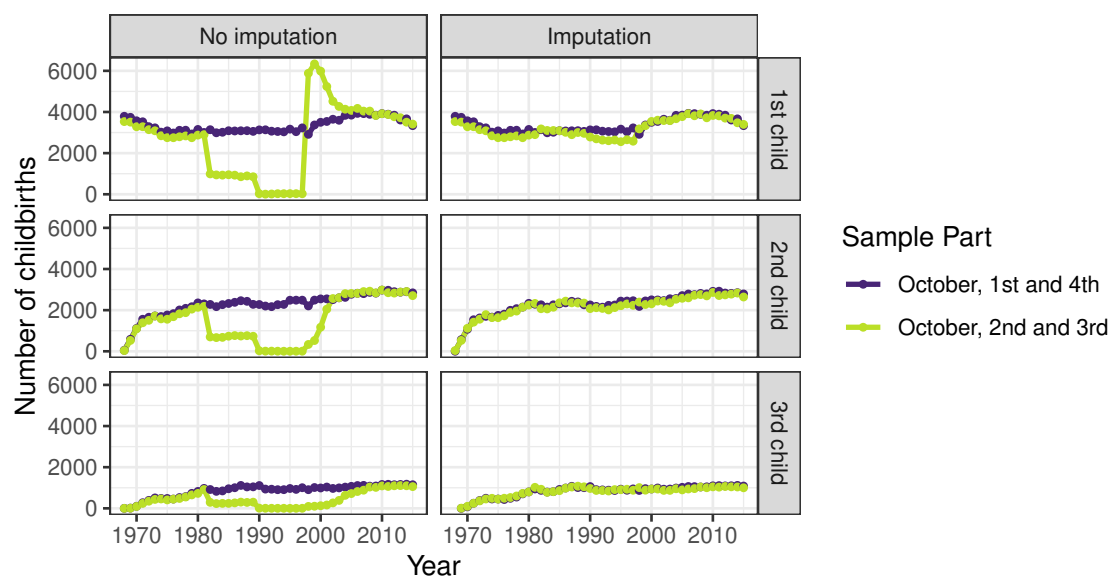
For each individual i belonging to the incomplete half of the sample, we impute first the year of first childbirth according to the following principles:

- if the first childbirth in birth records occurs before 1982, we regard it as the first childbirth;
- else,
 - if the earliest of years of childbirth she declared in the 1990 census is after 1982, we consider the earliest of these years and the year of the first childbirth as it appears in birth records as the year of the first childbirth;
 - else,
 - * if the earliest of years of childbirth she declared in the 1999 census is after 1982, we consider the earliest of these years and the year of the first childbirth as it appears in birth records as the year of the first childbirth;
 - * else,

- if birth records indicate that she has children, we consider the year of the first childbirth in birth records as the year of first childbirth;
- else, we assume that she has no children.

We then consider the n th childbirth with $n > 1$ as the minimum of years of childbirth within both birth records and censuses among years of birth that follow the computed year of the $n - 1$ th childbirth.

This approach does not take multiple births into account; more generally, it does not account for individuals who experience more than one childbirth per year. Despite this *caveat*, our approach matches the historical pattern in the complete half of the sample quite well. Figure 1.B.1 plots the number of childbirtths by rank of childbirth for each year since 1968 for both parts of the sample, relying on birth records only (left panel) and on our approach (right panel). While we still slightly underestimate first childbirtths that occur in the beginning of the 1980s or in the late 1990s in the incomplete half of the sample, our approach matches reasonably well the patterns observed in the complete half of the sample.

Figure 1.B.1 – Imputation of childbirths for individuals born October 2 and 3

1.C Age-Period-Cohort models

The major challenge in the simultaneous identification of λ , μ and ν stems from collinearity between age, cohort and period: age is equal to the current period less the year of birth. Several solutions have been explored in the sociological literature; e.g., [Mason et al. \(1973\)](#) assume that any two ages, periods or cohorts have the same effect, in addition to removing one dummy in each dimension. [Deaton and Paxson \(1994\)](#) and [Deaton \(1997\)](#) suggest a transformation of period effects to meet two requirements: (i) that time effects sum to zero, and (ii) that 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. Hence, the parameters of the model (λ, μ, ν) are identified, provided that $\lambda_c = 0$ and $\sum_{t=1}^T \nu_t(t-1) = 0$.

Because we rely on a sample that only contains individuals born on even-numbered years, we have to impose one additional normalization (see [Pora and Wilner, 2020](#), on this matter). Indeed, without additional normalizations the model is underidentified: during even-numbered years, individuals could face a systematic shock that is exactly offset by the fact that their ages are also even-numbered. Specifically, the two models:

$$\tilde{y}_{it} = \sum_c \lambda_c \mathbb{1}_{cohort_i=c} + \sum_a \mu_a \mathbb{1}_{age_{it}=a} + \sum_j \nu_j \mathbb{1}_{t=j} + \epsilon_{it}$$

and

$$\tilde{y}_{it} = \sum_c \lambda_c \mathbb{1}_{cohort_i=c} + \sum_a (\mu_a \mathbb{1}_{age_{it}=a} + \xi \mathbb{1}_{age_{it} \equiv 0[2]}) + \sum_j (\gamma_j \mathbb{1}_{t=j} - \xi \mathbb{1}_{t \equiv 0[2]}) + \epsilon_{it}$$

are observationally equivalent regardless of the identification of cohort coefficients.

This limitation leads us to further impose that both odd-year and even-year time effects sum to zero, i.e., to consider two restrictions: $\sum_{t>0} \nu_t = 0$ where $t = 2j$ and $\sum_t \nu_t = 0$ where $t = 2j + 1$. The corresponding transformation of time dummies $d_T = \mathbb{1}_{t=T}$ is written as follows:

$$d_t^* = \begin{cases} d_t - [\frac{t}{2}d_3 - \frac{t-2}{2}d_1] & t = 2j, j > 1 \\ d_t - [\frac{t-3}{2}d_3 + d_2 - \frac{t-3}{2}d_1] & t = 2j + 1, j > 1 \end{cases} \quad (1.13)$$

where $d_1^* = d_2^* = d_3^* = 0$. 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^* , d_2^* and d_3^* in the regression.

1.D Accounting decomposition

The log-change in total labor earnings between time $t - 1$ and time $t + k$ is written as

$$\Delta y_{t+k} = \log(\mathbb{E}[y_{i,t+k}]) - \log(\mathbb{E}[y_{i,t-1}]) \quad (1.14)$$

Δy_{t+k} can also be rewritten as

$$\Delta y_{t+k} = \log\left(\frac{\mathbb{E}\left[\frac{y_{i,t+k}}{y_{i,t-1}} y_{i,t-1}\right]}{\mathbb{E}[y_{i,t-1}]}\right) \quad (1.15)$$

This formulation is particularly relevant here since we require that all individuals be employed at time $t - 1$, so $y_{i,t-1} > 0$. Hence Δy_{t+k} is simply the log-average of individual changes $y_{i,t+k}/y_{i,t-1}$ weighted by initial earnings $y_{i,t-1}$.

Next, we use an accounting decomposition of labor earnings at the individual level. First, using the law of iterated expectations yields

$$\begin{aligned} \mathbb{E}[y_{i,t+k}] &= \mathbb{P}(d_{i,t+k} = 0) \mathbb{E}[y_{i,t+k} | d_{i,t+k} = 0] \\ &\quad + \mathbb{P}(d_{i,t+k} = 1) \mathbb{E}[y_{i,t+k} | d_{i,t+k} = 1] \end{aligned} \quad (1.16)$$

Since $d_{i,t+k} = 0 \Rightarrow y_{i,t+k} = 0$, the first term vanishes:

$$\begin{aligned} \Delta y_{t+k} &= \log(\mathbb{P}(d_{i,t+k} = 1)) + \log(\mathbb{E}[y_{i,t+k} | d_{i,t+k} = 1]) - \log(\mathbb{E}[y_{i,t-1}]) \\ &= \underbrace{\log(\mathbb{P}(d_{i,t+k} = 1))}_{\text{Participation}} + \underbrace{\log(\mathbb{E}[y_{i,t-1} | d_{i,t+k} = 1]) - \log(\mathbb{E}[y_{i,t-1}])}_{\text{Selection}} \\ &\quad + \underbrace{\log(\mathbb{E}[y_{i,t+k} | d_{i,t+k} = 1]) - \log(\mathbb{E}[y_{i,t-1} | d_{i,t+k} = 1])}_{\Delta y_{t+k}^{\text{Participants}}} \end{aligned} \quad (1.17)$$

We are thus left with the decomposition of the latter term $\Delta y_{t+k}^{\text{Participants}}$; for these participants, all components of labor earnings – days, hours and hourly wages – are

observed in the data. Then,

$$\begin{aligned} \Delta y_{t+k}^{\text{Participants}} &= \log \left(\underbrace{\frac{\mathbb{E} \left[\frac{w_{i,t+k}}{w_{i,t-1}} x_{i,t+k} h_{i,t+k} w_{i,t-1} \mid d_{i,t+k} = 1 \right]}{\mathbb{E} [x_{i,t+k} h_{i,t+k} w_{i,t-1} \mid d_{i,t+k} = 1]}}_{\text{Hourly wages growth}} \right) \\ &\quad + \log \left(\frac{\mathbb{E} [x_{i,t+k} h_{i,t+k} w_{i,t-1} \mid d_{i,t+k} = 1]}{\mathbb{E} [x_{i,t-1} h_{i,t-1} w_{i,t-1} \mid d_{i,t+k} = 1]} \right) \end{aligned} \quad (1.18)$$

We continue to perform similar substitutions in the second term with respect to the two remaining components (hours and days). It follows that

$$\begin{aligned} \underbrace{\Delta y_{t+k}}_{\text{Labor earnings changes}} &= \underbrace{\log (\mathbb{P}(d_{i,t+k} = 1))}_{\text{Participation}} \\ &\quad + \underbrace{\log \left(\frac{\mathbb{E} [y_{i,t-1} \mid d_{i,t+k} = 1]}{\mathbb{E} [y_{i,t-1}]} \right)}_{\text{Selection}} \\ &\quad + \underbrace{\log \left(\frac{\mathbb{E} \left[\frac{x_{i,t+k}}{x_{i,t-1}} x_{i,t-1} h_{i,t-1} w_{i,t-1} \mid d_{i,t+k} = 1 \right]}{\mathbb{E} [x_{i,t-1} h_{i,t-1} w_{i,t-1} \mid d_{i,t+k} = 1]} \right)}_{\text{Changes in Days Worked}} \\ &\quad + \underbrace{\log \left(\frac{\mathbb{E} \left[\frac{h_{i,t+k}}{h_{i,t-1}} x_{i,t+k} h_{i,t-1} w_{i,t-1} \mid d_{i,t+k} = 1 \right]}{\mathbb{E} [x_{i,t+k} h_{i,t-1} w_{i,t-1} \mid d_{i,t+k} = 1]} \right)}_{\text{Changes in Hours Per Day}} \\ &\quad + \underbrace{\log \left(\frac{\mathbb{E} \left[\frac{w_{i,t+k}}{w_{i,t-1}} x_{i,t+k} h_{i,t+k} w_{i,t-1} \mid d_{i,t+k} = 1 \right]}{\mathbb{E} [x_{i,t+k} h_{i,t+k} w_{i,t-1} \mid d_{i,t+k} = 1]} \right)}_{\text{Hourly Wage Growth}} \end{aligned} \quad (1.19)$$

This accounting identity clarifies that the (reweighted) log-average of individual earnings' changes can be decomposed into the sum of (reweighted) log-average of individual changes for each component, and a selection term.

1.E Right-censoring and measurement error

Our definition of control and treatment groups, despite being practical, raises some issues. First, due to right-censoring, individuals in our control group are not of the same age as those in our treatment group. Second, and for the same reason, our treatment effect estimate corresponds to the difference in labor market outcomes between parents with k children and individuals with $k - 1$ children over the lifetime; this is true for old cohorts, but our estimate for younger cohorts might be spuriously affected by selection bias, namely, differences between parents of k children who experience childbirths quite early and parents who eventually have k children but do so later in life. Third, the definition of our treatment as experiencing the k th childbirth during year t might be blurred by the timing of labor supply decisions mainly because women are entitled a maternity leave that begins several weeks before childbirth and ends several months after. Choosing year $t - 1$ as a reference for labor market outcomes may therefore lead to biases with respect to childbirths that occur in the very beginning of the year since part of the childbirth effect might already have happened.

We address all three issues by providing several estimations based on alternative definitions of control and/or treatment groups:

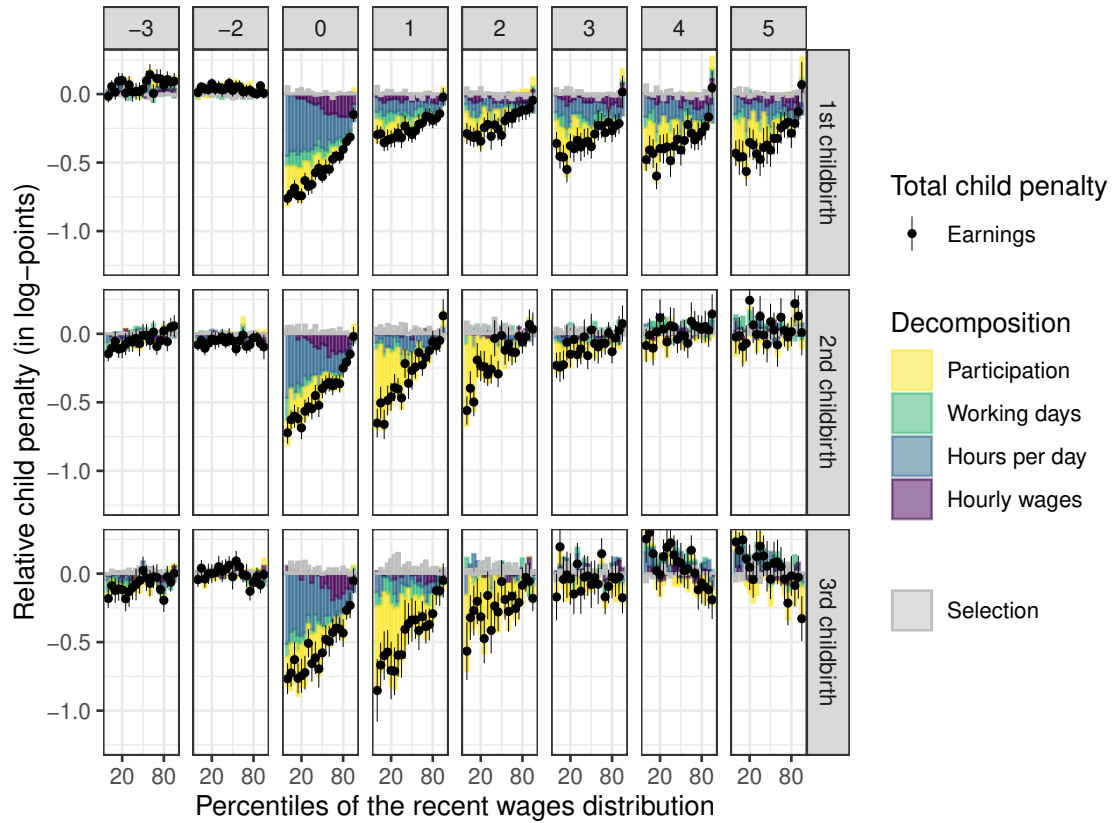
1. We define our n th control group as individuals that experience $n - 1$ childbirths according to the data, as of age randomly drawn from the empirical distribution of age at the n th childbirth within education \times cohort cells. This allows us to assess robustness with respect to the age difference between control and treatment groups (see Figures 1.E.1 and 1.E.2).
2. We restrict our analysis to individuals born in 1975 or before: such individuals are most likely to have made all of their fertility decisions by year 2015 (see Figures 1.E.3 and 1.E.4).
3. We define our n th control group as individuals who have n children according to the data as of time t and do not experience any childbirths between $t - 1$ and $t + k$ (see Figures 1.E.5 and 1.E.6). This strategy is closer to that of [Kleven et al. \(2019b\)](#) in that it relies on the timing of the n th childbirth among those who indeed have n children.
4. We restrict our n th treatment group to individuals who experience the n th

childbirth during the second quarter, i.e., between April and June: their maternity leaves do not begin before January and do not end after December (see Figures 1.E.7 and 1.E.8).

Our findings prove robust to these alternative definitions.


Additionally, our measure of the causal impact of childbirth rests on the assumption that treatment effects are time-homogeneous, i.e., that childbirths occurring in 1998 have the same causal impact as those that occurred in 2015 if they are considered after the passage of the same amount of time. This assumption justifies reliance of our estimates on a time-varying window: the impact of childbirth at time t is estimated on all childbirths from 1998 to 2015, while our estimate of the impact of childbirths at time $t + 5$ only relies on childbirths that occurred before 2011. The credibility of this assumption might be problematic given that the institutional background varied over the time period, e.g. the PAJE reform took place in 2004, and another change in parental leave rules, which was a slight change in the incentives to split parental leave between parents, happened in 2015. Nevertheless, we replicate our analysis while restricting it to childbirths during 2000-2010; hence this compositional change does not distort our estimates of dynamic treatment effects: all treatment effects for all durations of time to childbirth are computed on the very same sample. Additionally, by choosing 2000 as the beginning of the estimation timespan, we can ignore the fact that the selection condition in our sample is harsher for childbirths in 1998-1999 due to left-censoring issues in the data. Figures 1.E.9 and 1.E.10 display our estimates and show that our approach is completely robust with respect to these concerns.

Figure 1.E.1 – Consequences of childbirth for women’s labor outcomes: a comparison with a control group determined at the imputed age of childbirth



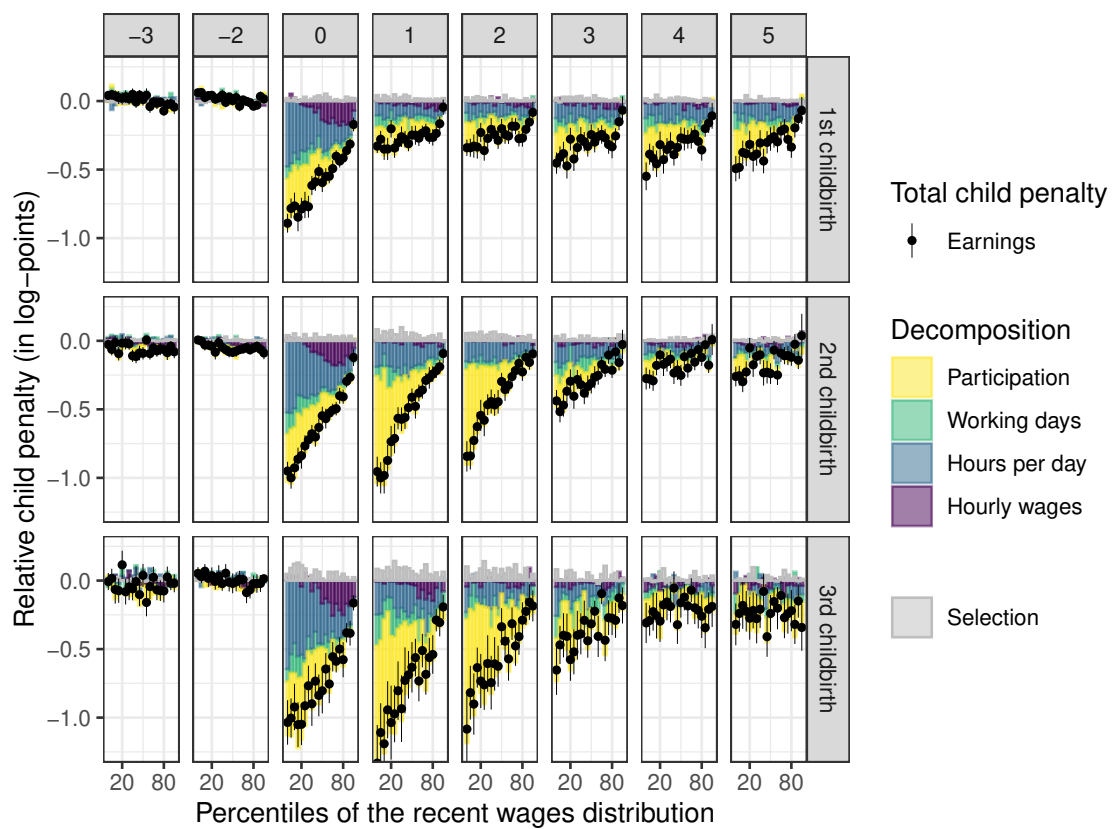
Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see 1.6) for various values of time-to-childbirth expressed in years. The control group is determined at age randomly drawn from the distribution of age at the n th childbirth within gender \times cohort \times education cells. Bootstrapped standard errors using 100 replications are clustered at the individual level.

Figure 1.E.2 – Consequences of childbirth for men’s labor outcomes: a comparison with a control group determined at the imputed age of childbirth



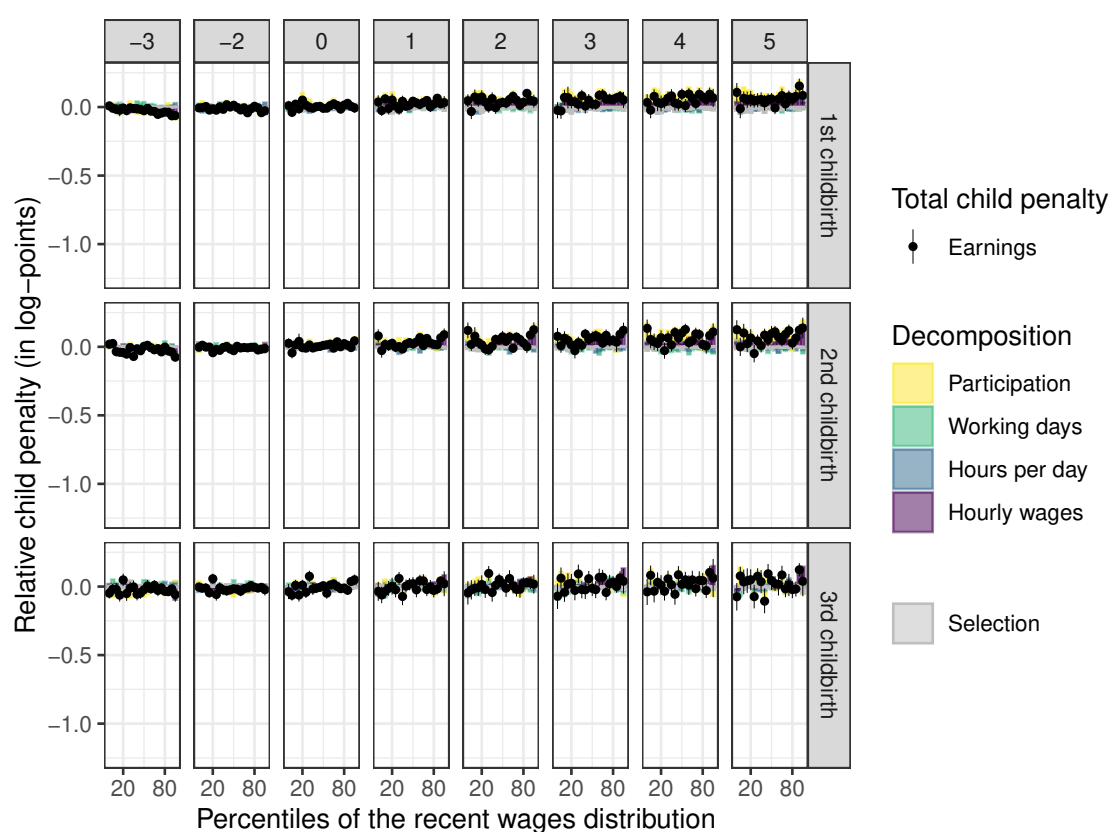
Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see 1.6) for various values of time-to-childbirth expressed in years. The control group is determined at age randomly drawn from the distribution of age at the n th childbirth within gender \times cohort \times education cells. Bootstrapped standard errors using 100 replications are clustered at the individual level.

Figure 1.E.3 – Consequences of childbirth for women’s labor outcomes: restriction to older cohorts that have made complete fertility decisions



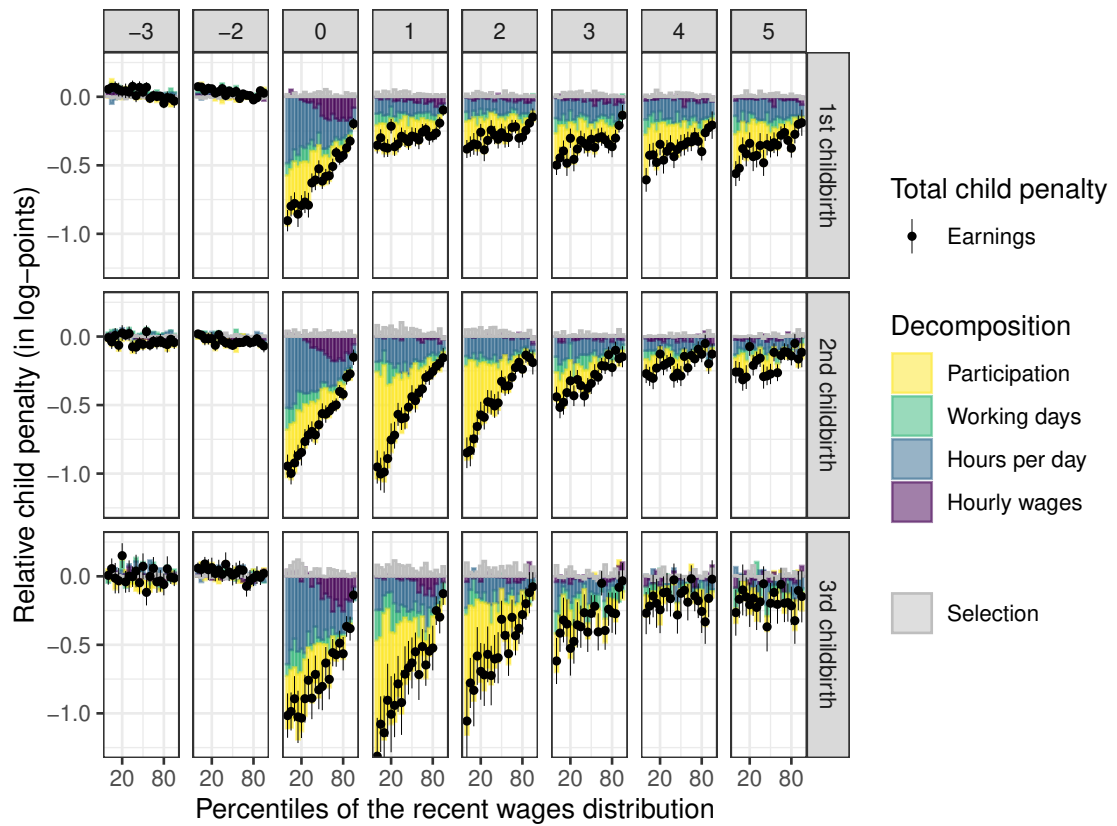
Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see 1.6) for various values of time-to-childbirth expressed in years. The sample is restricted to individuals born in 1975 or earlier. Bootstrapped standard errors using 100 replications are clustered at the individual level.

Figure 1.E.4 – Consequences of childbirth for men’s labor outcomes: restriction to older cohorts that have made complete fertility decisions




Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see 1.6) for various values of time-to-childbirth expressed in years. The sample is restricted to individuals born in 1975 or earlier. Bootstrapped standard errors using 100 replications are clustered at the individual level.

Figure 1.E.5 – Consequences of childbirth for women’s labor outcomes: identification based on the timing of the k th childbirth



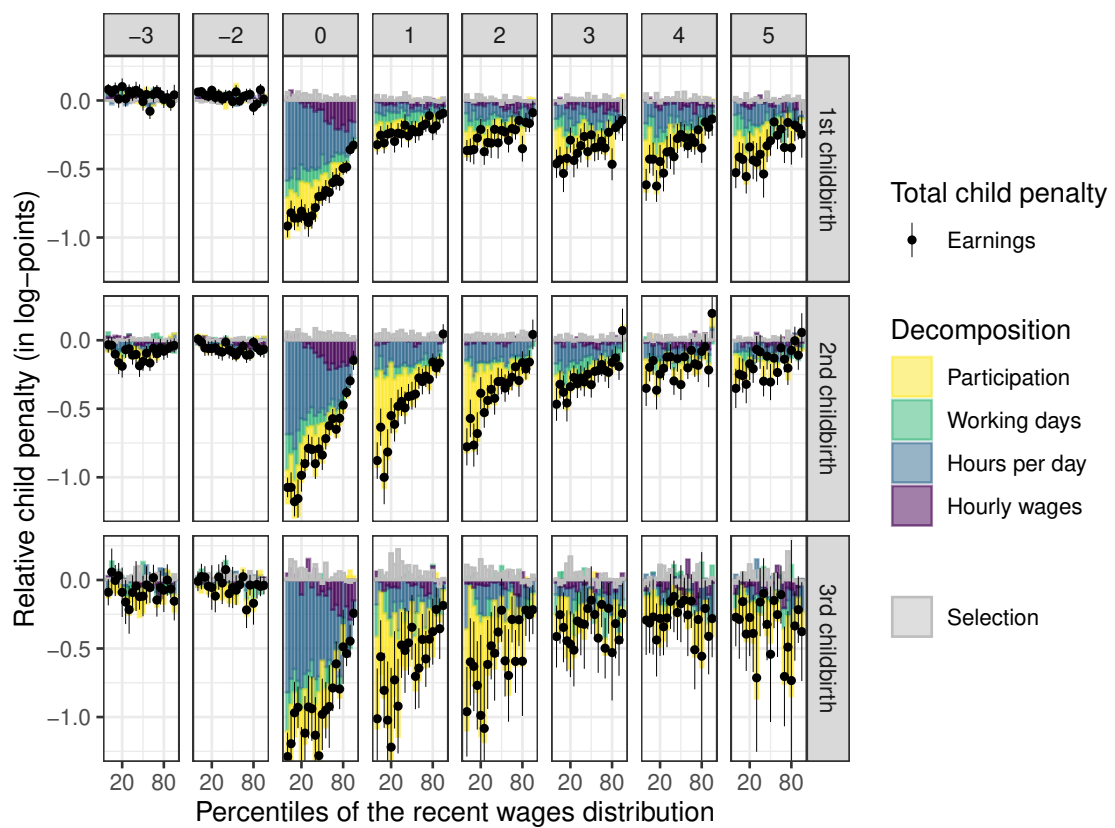
Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see 1.6) for various values of time-to-childbirth expressed in years. The control group includes individuals with exactly $n - 1$ children in 2015 and who do not experience childbirth between $t - 1$ and $t + k$. Bootstrapped standard errors using 100 replications are clustered at the individual level.

Figure 1.E.6 – Consequences of childbirth for men’s labor outcomes: identification based on the timing of the k th childbirth



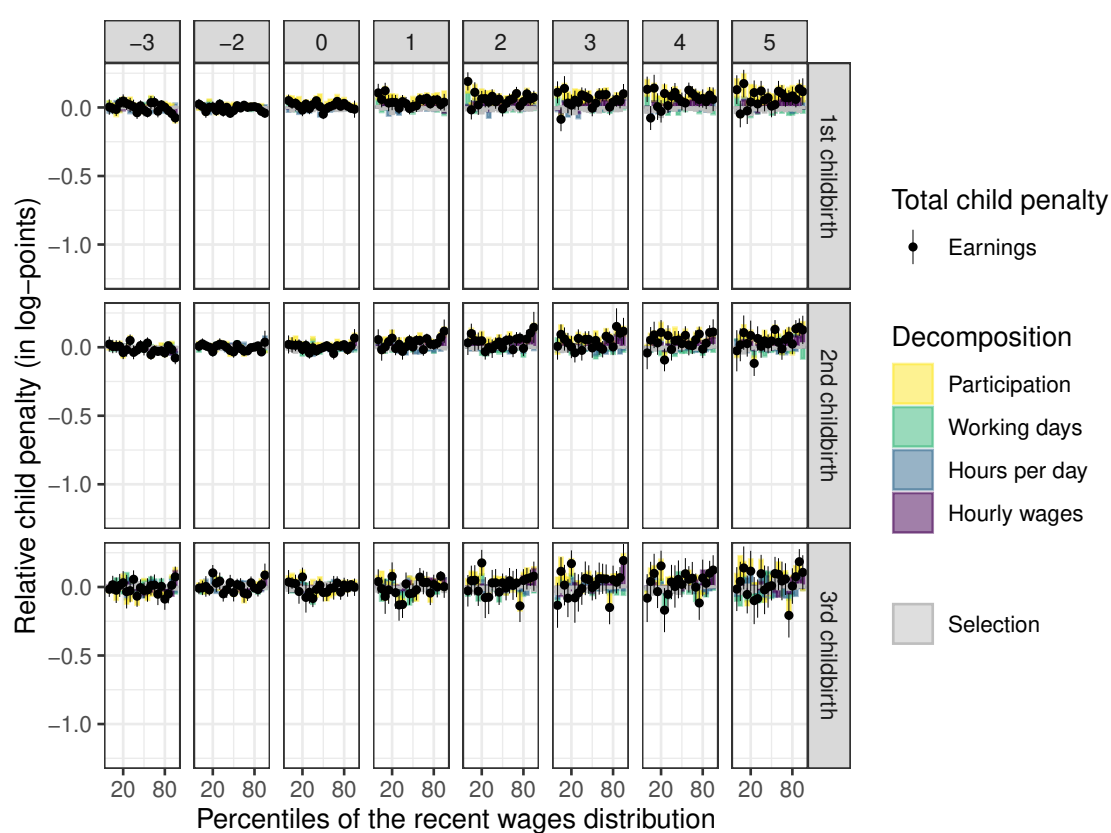
Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see 1.6) for various values of time-to-childbirth expressed in years. The control group includes individuals with exactly n children in 2015 and who do not experience childbirth between $t - 1$ and $t + k$. Bootstrapped standard errors using 100 replications are clustered at the individual level.

Figure 1.E.7 – Consequences of childbirth for women’s labor outcomes: restriction to childbirths in the second quarter



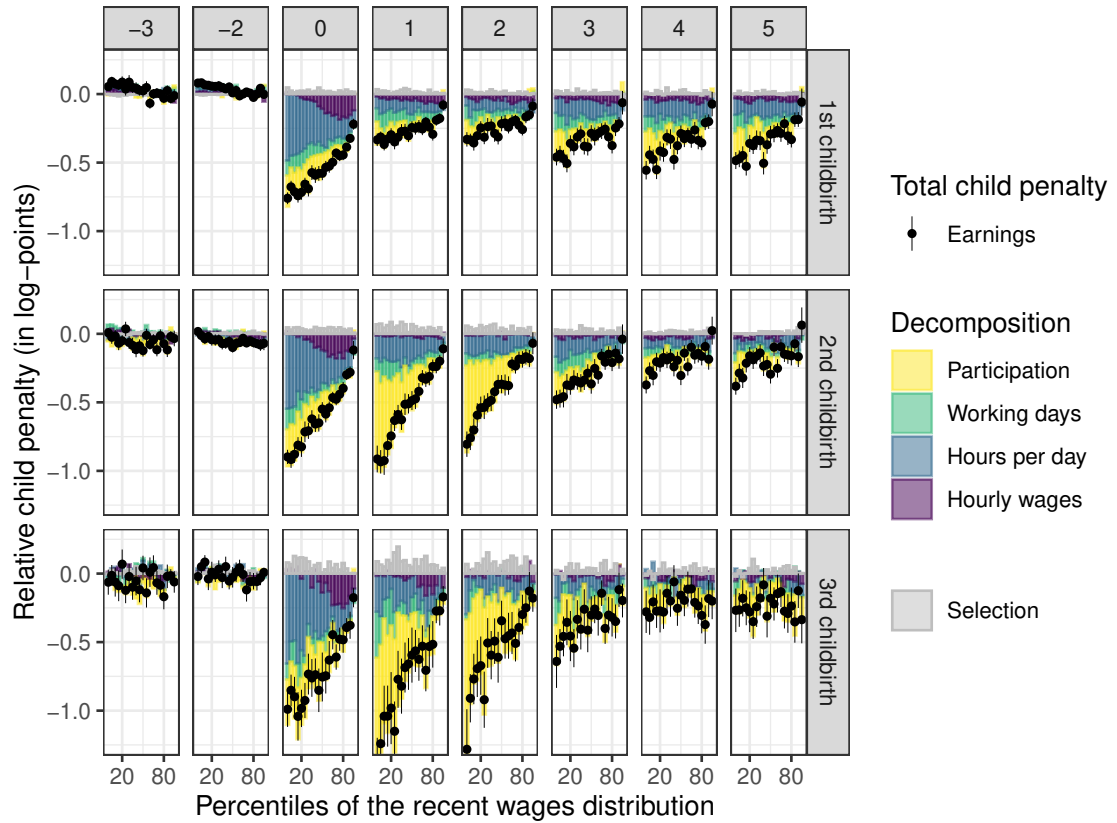
Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see 1.6) for various values of time-to-childbirth expressed in years. The treated group is restricted to individuals that experience the n th childbirth during the second quarter of year t . Bootstrapped standard errors using 100 replications are clustered at the individual level.

Figure 1.E.8 – Consequences of childbirth for men’s labor outcomes: restriction to childbirths in the second quarter




Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see 1.6) for various values of time-to-childbirth expressed in years. The treated group is restricted to individuals that experience the n th childbirth during the second quarter of year t . Bootstrapped standard errors using 100 replications are clustered at the individual level.

Figure 1.E.9 – Consequences of childbirth for women’s labor outcomes: restriction to childbirths in 2000-2010



Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see 1.6) for various values of time-to-childbirth expressed in years. The treatment time is restricted to years 2000 to 2010. Bootstrapped standard errors using 100 replications are clustered at the individual level.


Figure 1.E.10 – Consequences of childbirth for men’s labor outcomes: restriction to childbirths in 2000-2010



Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see 1.6) for various values of time-to-childbirth expressed in years. The treatment time is restricted to years 2000 to 2010. Bootstrapped standard errors using 100 replications are clustered at the individual level.

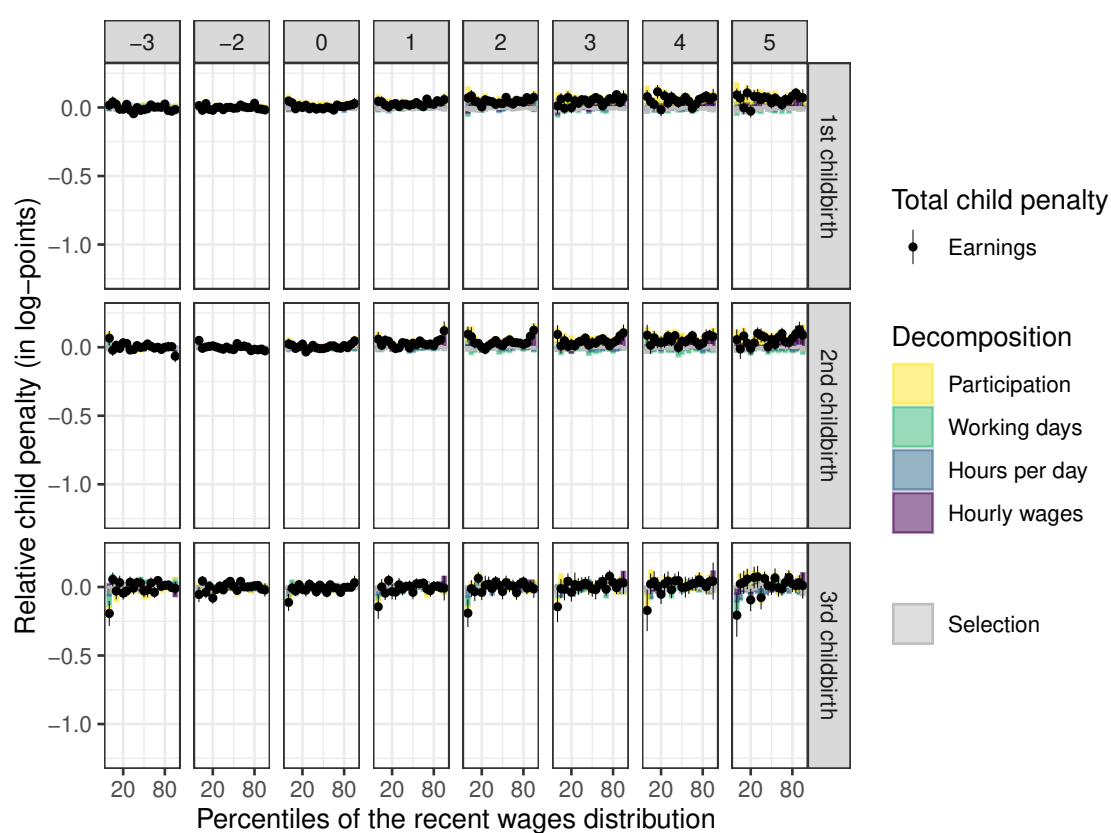
1.F Alternate ranking

Figure 1.F.1 – Consequences of childbirth for women’s labor outcomes: alternate ranking



Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see 1.6) for various values of time-to-childbirth expressed in years. The ranking is based on average normalized hourly wages as measured between $t - 7$ and $t - 3$. Bootstrapped standard errors using 100 replications are clustered at the individual level.

Figure 1.F.2 – Consequences of childbirth for men’s labor outcomes: alternate ranking

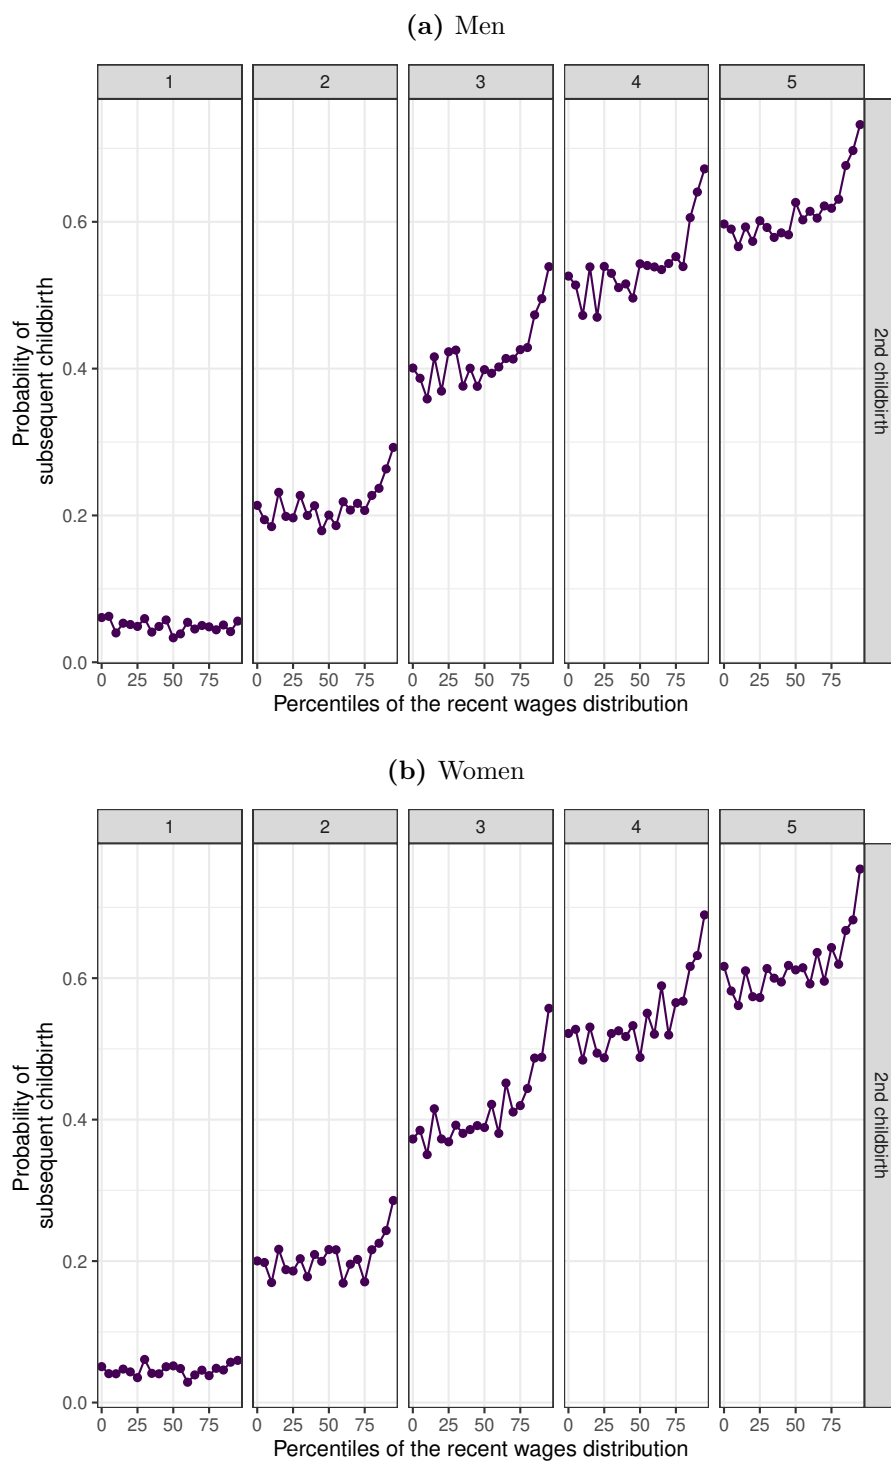


Each panel displays the estimates of child penalties obtained by the difference-in-difference method (see 1.6) for various values of time-to-childbirth expressed in years. The ranking is based on average normalized hourly wages as measured between $t - 7$ and $t - 3$. Bootstrapped standard errors using 100 replications are clustered at the individual level.

1.G Subsequent fertility decisions

By construction, our estimates of child penalties aggregate the consequences of a given childbirth with all subsequent fertility decisions. To check that the heterogeneity in child penalties along the wage distribution is not driven by high-achieving mothers being more prone to restrict their future fertility, we compute the probability of having at least a second child a few years after the first childbirth. Figure [1.G.1](#) displays the corresponding estimates.

For both men and women, this probability is a non-decreasing function of pre-childbirth hourly wages. In other words, up to 5 years after the first childbirth, parents at the top of the pre-childbirth wage distribution are, if anything, less likely than their counterparts from the bottom of the distribution to restrict themselves to one child only. Hence our results on the heterogeneity of child penalties are not driven by mothers with high hourly wages having less children.

Figure 1.G.1 – Probability of having a 2nd child along the wage distribution

Each panel displays the estimates of the probability of having a second child for various values of time-to-first-childbirth expressed in years.

1.H Childcare preferences

We rely on the complemented LFS data to estimate the relation between preferences towards childcare and hourly wages. To this end, we estimate a log-hourly wage equation in which we introduce a dummy variable which indicates that the individual views childcare provided by parents as the best childcare solution. Table 1.H.1 displays our results. Consistent with Figure 1.10, we find a small negative correlation between preferences and hourly wages among those for which information regarding hourly wages is available in the LFS, i.e. those who were in salaried employment at the time they were surveyed. This correlation is not significant at usual levels so that we cannot reject the null hypothesis that childcare preferences are not correlated with hourly wages in this selected sample. Additionally, this correlation reverses for men, and shrinks substantially for women when education is accounted for.

Relying on observed hourly wages only, this result may yet be misleading due selection bias. Specifically, women with the lowest potential hourly wages and the largest taste for childcare are more likely to leave the labor force, which may lead to understate the correlation between childcare preferences and potential hourly wages. We therefore estimate a Heckman sample selection model. We take advantage of one question of the survey that asked respondent who were either out of the labor force or worked part-time whether their labor supply decision was the result of problems regarding childcare. We therefore rely on this dummy variable, assuming that it is equal to 0 for all individuals working full-time; the exclusion restriction is that childcare issues affect the participation decision but do not directly affect hourly wages. It turns out that women who hold the most conservative views regarding childcare or who face problems regarding the access to external childcare are more likely to be outside the labor force. However, correcting for sample selection does not change substantially our estimates of the relation between hourly wages and childcare preferences. As a result, we do not find significant effects of childcare preferences on potential hourly wages.

Table 1.H.1 – Correlation between childcare preferences and hourly wages

	(1)	(2)	(3)	(4)
Log hourly wage				
Parents	-0.055 (0.035)	0.011 (0.031)	-0.055 (0.035)	0.010 (0.031)
Parents \times Female	-0.032 (0.058)	-0.037 (0.049)	-0.030 (0.061)	-0.033 (0.049)
Selection				
Parents			0.16 (0.098)	0.18 (0.098)
Parents \times Female			-0.52 (0.14)	-0.52 (0.14)
Childcare issues			0.15 (0.41)	0.075 (0.38)
Childcare issues \times Female			-1.00 (0.43)	-0.92 (0.41)
Correlation between residuals				
$\hat{\rho}$			-0.015 (0.15)	-0.034 (0.089)
Additional controls				
Age youngest child	✓	✓	✓	✓
Age	✓	✓	✓	✓
Education		✓		✓
Observations	1151	1151	2550	2550

The sample is restricted to individuals with children aged 3 or less. All covariates are further interacted with a female dummy. Standard errors in parentheses.

Chapter 2

**Keep working and spend less?
Collective childcare and parental
earnings in France**

2.1 Introduction

In international comparisons, France is widely seen as a success in terms of family policies that promote the work-family balance and gender equality. Among both OECD and EU countries, it ranks high in terms of fertility rate, female employment rate and formal childcare coverage (see e.g. [OECD, 2011](#)). In contrast to other countries, childcare arrangements in France are extremely diverse: its long-lasting institutional history has led to the coexistence of paid parental leave and highly subsidized formal childcare services, the latter including a continuum from individual at-home childcare to collective services provided by daycare centers. This unique diversity is supported by policy-makers and the general public, and is assumed to provide families with freedom to choose the childcare arrangements most suited to their heterogeneous preferences and constraints.

In this paper, I highlight a consequence of this institutional setting. Namely, this diversity leaves room for potentially large substitution effects across childcare solutions. Consequently, massive investment plans aimed at increasing the overall provision of formal childcare may simply crowd out other subsidized solutions, instead of further enhancing the work-life balance and gender equality. Specifically, I investigate the consequences in terms of both parental labor earnings and labor supply, and childcare choices, of a series of national plans launched in the 2000s and designed to increase the supply of daycare centers to provide particularly affordable collective childcare for very young children. My main results show that: (i) these plans did not trigger any substantial change in the parental labor supply, and especially that of mothers; (ii) instead, they resulted in families shifting away from more expensive individualized childcare solutions.

While focused on the unique French setting, this paper is relevant to more general questions regarding the impact of affordable childcare on maternal labor supply. Indeed, null effects have sometimes been reported in the literature (e.g. [Fitzpatrick, 2010](#); [Havnes and Mogstad, 2011a](#)), which has been usually attributed to substitution effects across childcare solutions. However, with few exceptions, this mechanism remains somewhat speculative because it usually involves the crowding-out of informal childcare solutions (e.g. childcare provided by a relative or a neighbor). The main problem is that these informal arrangements are not observed in the data with sufficient precision and frequency to correctly identify the causal effect of affordable childcare supply on childcare choices when these policies fail to enhance mothers' la-

bor outcomes. Informal primary childcare providers are quite uncommon in France, however, which implies that gathering data on the multiple formal childcare solutions is generally sufficient to cover almost all relevant childcare choices. I am therefore able to provide clear evidence that the null effects of affordable childcare provision on maternal labor supply do indeed arise from crowding-out effects.

My empirical approach focuses on the staggered expansion of affordable collective childcare across narrow geographical areas in response to a succession of national plans that aimed at increasing the overall collective formal childcare provision. Specifically, I leverage differences in the timing of major expansion events *across* municipalities, *within* groups of municipalities that experienced increases of similar magnitude, to identify the causal effect of affordable childcare on parents' labor earnings, labor supply and childcare choices. I apply this framework to a combination of detailed administrative datasets: childcare and parental leave records kept by the Family branch of the French Social security, as well as both cross-sectional and longitudinal birth records and payroll tax data.

I find that these sharp increases in affordable collective childcare provision at the municipal level did not trigger any substantial change in parental labor outcomes. Specifically, my estimates are incompatible with causal effects of childcare expansions on maternal employment larger than 0.05 percentage points per percentage point increase in the childcare coverage rate

I then shed light on the underlying mechanisms that generate these null effects. Firstly, I consider the possible substitution with paid parental leave to which most parents of very young children are entitled. While most empirical studies consider childcare provision and parental leave as two separate policies, this parameter bears substantial policy relevance as it predicts whether a change in the childcare provision is likely to affect the demand for parental leave or not. Consistent with my labor supply estimates, I find that the expansion of affordable collective childcare does not trigger any change in the take-up of parental leave benefits, which suggests that these substitution effects are limited at best.

Secondly, I focus on the supply of other formal and more costly childcare solutions, i.e. childminders and nannies providing at-home childcare. Applying the same approach, I provide evidence of a very substantial crowding-out of these childcare solutions after collective childcare expansions. Specifically, in municipalities with the sharpest increases in affordable collective childcare provision, the medium-run drop in individualized childcare supply is equivalent in magnitude to that of the

increase in collective childcare provision. This implies that the increased childcare capacity of daycare centers likely benefits parents who would have otherwise turned to individualized and more expensive formal childcare solutions.

This suggests that these families have a high propensity to rely on formal childcare solution, regardless of the availability of affordable childcare, which may stem from either a strong taste for working mothers, or strong incentives for mothers to remain in the labor force (for instance if the hourly price of individualized childcare is only a fraction of mothers' hourly wages). It does *not* follow that, due to either preferences or incentives, the numerous families who did not benefit from a collective childcare place would not change their labor supply decision in response to childcare places being made available to them. Indeed, my estimates are only informative about the subpopulation of families who are offered a childcare place thanks to the local childcare expansion, but would not have been so before the expansion took place. Extrapolating these effects to never-treated families is not straightforward, and would likely require additional data on the allocation of childcare places, at both the application and the selection level of the process.

Literature When seeking to identify the labor supply effects of childcare provision the main empirical challenge to overcome is the fact that childcare and labor supply decisions are made jointly: the causal impact of childcare on labor supply cannot be identified from the correlation between actual childcare and labor supply choices. As a result, researchers have resorted to either a careful specification of the joint decision process (e.g. Heckman, 1974; Michalopoulos et al., 1992; Domeij, 2013; Bick, 2016) or quasi-experimental evidence arising from plausibly exogenous policy changes (e.g. Gelbach, 2002; Baker et al., 2008; Fitzpatrick, 2010; Bauernschuster and Schlotter, 2015; Gathmann and Sass, 2018; Carta and Rizzica, 2018).

Especially relevant to this paper are studies that infer the causal impact of affordable childcare on maternal labor outcomes by exploiting heterogeneity between geographical areas in the timing of publicly subsidized childcare expansions in response to national-level policy reforms (Berlinski and Galiani, 2007; Havnes and Mogstad, 2011a; Nollenberger and Rodríguez-Planas, 2015; Yamaguchi et al., 2018; Andresen and Havnes, 2019; Müller and Wrohlich, 2020). Broadly speaking, such papers manage to get around the endogeneity with respect to labor supply of both individual childcare choices and the local childcare availability by relying on a fuzzy difference-in-difference framework akin to that of Duflo (2001). Specifically, they

leverage the fact that some areas experience large and sudden increases in affordable childcare provision, while other do not, or may experience them later on. The former are thus considered as a treated group, while the latter are used like a control group, under the assumption that, absent the treatment, labor outcomes in the treated group would have evolved in the same way as those in the control group, so as to capture any change that occurs at a national level. My identification strategy relies on a variation of this approach.

In terms of results, this literature is somewhat contrasted between papers that find substantial positive effects of affordable childcare provision on maternal labor supply, and others that emphasize null effects. In the US and Canada, [Blau and Currie \(2006\)](#) report estimates of maternal labor supply elasticity with respect to the price of childcare. Across the 20 studies analyzed, these estimates vary from -3.60 to +0.06. For a more recent perspective on the literature, [Morrissey \(2017\)](#) reports elasticities that range from -1.1 to -0.025 in the US. Variation may stem from the age of the targeted children, the educational attainment or labor force attachment of the targeted mothers, or broader variation in national or historical context; even so, the results are not always easy to reconcile. When combined with quasi-experimental approaches to child penalties such as that of [Kleven et al. \(2019\)](#), the difference-in-difference approach of [Havnes and Mogstad \(2011a\)](#) yields contrasting results: [Nix and Andresen \(2019\)](#) suggest that early childcare has the potential to alleviate the child penalty in Norway, whereas [Kleven et al. \(2020\)](#) emphasize that increased childcare provision has no effect on the child penalty in Austria.

Null effects are thus not uncommon in this literature, and have been attributed to substitution across childcare solutions. To date, [Cascio \(2009\)](#) provides the most compelling evidence as to these crowding-out effects, but empirical facts regarding such effects remain otherwise scarce. While [Baker et al. \(2008\)](#) provide direct evidence of crowding-out effects, although in a context where maternal labor supply effects are actually positive, [Asai et al. \(2015\)](#) suggest that these effects may explain the observed heterogeneity in the maternal labor supply effect between two-generation and three-generation families, in a context in where childcare is frequently provided by grandparents. [Bassok et al. \(2014\)](#) document substitution effects between public and private childcare, with the magnitude of crowding-out depending on the type of intervention (e.g. a voucher program as opposed to direct public-sector childcare provision), but do not provide evidence as to the labor supply consequences of such crowding-out effects. These substitution effects are also relevant

to more general questions about the impact of regulation on the childcare market (Hotz and Xiao, 2011).

Few researchers have examined the French setting. Among them, both Choné et al. (2004) and Allègre et al. (2015) use a joint model of childcare choices and labor supply decisions, but reach different conclusions as to the effect of childcare prices on childcare choices and maternal labor supply. This may arise from differences in the level of detail of the childcare data they use. Closer to a quasi-experimental approach, Maurin and Roy (2008) examine the difference between families that obtained a childcare place and those who did not among all families who applied in a particular city, and find a positive effect on maternal labor supply. Goux and Maurin (2010) focus on the availability of pre-school places for 2-years olds, and find a positive impact for single mothers, but not for mothers with a cohabiting partner. Lastly, Givord and Marbot (2015) examine the effects of a policy reform implemented in 2004 that led to a sharp decrease in childcare costs for some families; they find a positive but small impact on maternal labor supply.

The remainder of the paper is organized as follows. The next section presents the institutional setting. Section 2.3 describes the data and section 2.4 details the identification strategy. Section 2.5 presents the results on parental earnings and labor supply. Section 2.6 investigates the underlying mechanisms, i.e. substitution across childcare solutions, and lastly, section 2.7 concludes.

2.2 Institutional setting

2.2.1 Early childcare coverage

France is among OECD countries with the broadest access to early childcare outside the home: in 2016, over 56% of children aged 2 or less were enrolled in early childcare, a share that only Denmark, Belgium and Iceland exceed (OECD, 2016). I focus exclusively on childcare for children under age 3 given that children in France can enter pre-school from age 3 and the enrollment rate is over 99%.

France has achieved this broad childcare coverage by fostering very diverse childcare arrangements, with daycare centers representing only a fraction of the total. Formal individualized childcare solutions, such as childminders and, to a lesser extent, individual at-home childcare are also quite common. Few parents rely heavily

on informal solutions in France: less than 3% of families with young children relied on a relative as their primary childcare provider in 2013 (Villaume and Legendre, 2014).

In this paper, I focus on one type of formal childcare provided outside the home, that of daycare centers, i.e. formal collective solutions, in contrast with formal individualized solutions (e.g. childminders or at-home childcare provided by nannies) or informal solutions (e.g. childcare provided by relatives) or lastly the decision not to rely on an external childcare provider. These collective solutions, coined as *Établissements d'Accueil du Jeune Enfant* (EAJE) accounted for 31% of total theoretical formal early childcare capacity in 2014 (IGAS/IGF, 2017).

2.2.2 EAJE-PSU facilities

Broadly speaking, EAJE facilities provide childcare to children up to age 6. However, because almost 100% of children attend school from age 3, they are more generally targeted towards children aged 0 to 2.¹ These facilities are often run by local authorities, sometimes through an association.

Specifically, I investigate the provision of childcare by EAJE facilities funded under the *Prestation de Service Unique* (PSU) scheme. Local offices (*Caisse d'Allocations Familiales*, CAF) of the Family branch (*Caisse Nationale d'Allocations Familiales*, CNAF) of the French Social Security system fund a large share of EAJE facilities through this scheme. To obtain this funding, it is required that an EAJE facility bases its pricing on a national fee schedule that makes it the cheapest formal childcare solution for families.² Figure 2.1 emphasizes this fact by displaying estimates of the prices paid by families across formal childcare solutions, and the corresponding burden for public finances.

Allocation of EAJE-PSU childcare places is decided at local level. Criteria may vary from one place to another, but they generally take into account the parents' place of residence, their employment status and the socio-economic background of the family. The only universal criterion is the municipality of residence (Onape, 2012).

¹Less than 1% of children aged 3 to 6 attend EAJE facilities in the evening (Villaume and Legendre, 2014).

²Appendix 2.A.1 details other criteria that EAJE facilities have to meet, and the pricing of childcare places.

2.2.3 National expansion plans

Until the early 2000s, the development of EAJE facilities was mostly decided by local authorities. In June 2000, the first national *plan crèche* (daycare center plan) was launched. Its main aim was to increase the availability of formal collective childcare, either by expanding pre-existing facilities, or by creating new ones. Since then, several other national plans have followed: the 9th *plan crèche* was launched in 2018. These plans are coordinated at national level by the CNAF, and implemented by local authorities with the help of local CAF offices. Local CAF offices usually allocate subsidies based on the number of formal childcare places relative to the number of children aged 3 or less, as observed at the municipality level.³

Between 2000 and 2016, 150 000 new subsidized childcare places were created, 2/3 of which were so through the opening of new facilities. Whether directly subsidized by these plans or not, the number of collective childcare places increased by 70 000 between 2007 and 2015, my period of interest. This is a relatively modest increase, at the national level, given that the number of children aged 2 or less over the same period was between 2.3 and 2.4 million.

2.2.4 Parental leave policies

Benefits may be granted when a parent interrupts his or her career or opts to work part-time (previously *Complément Libre Choix d'Activité* (CLCA) and now *Prestation Partagée d'Éducation de l'enfant* (PreParE)). Additionally, parents are entitled by law to extend the duration of their parental leave if they are not offered a formal childcare place.⁴ This policy is effective with the first birth and provides a fixed non-means-tested monthly amount for the maximum duration of 6 months; the duration increases up to 2 years from the second child on.⁵ Contrary to Sweden, for instance, the benefits do not depend on parents' past income: they amount to approximately €400 per month in the case of career interruption and to nearly €200 in the case of 80% part-time work.

Lastly, in Appendix 2.A.4, I analyze survey data to determine parents' childcare preferences, and the constraints affecting their childcare choices. The main lessons are that (i) there is strong parental demand for collective childcare; (ii) this demand

³See Appendix 2.A.2 for additional details

⁴Article L531-4 of the French Social Security Code.

⁵Appendix 2.A.3 further details this policy.

is likely not met by current supply thereof and; (iii) the lack of such collective childcare solutions may have negative consequences for maternal labor supply.

2.3 Data

My analysis combines several administrative records to recover (i) a measure of the supply of formal and collective childcare at a narrow geographical level and; (ii) labor market trajectories and fertility decisions of a large sample of individuals of whom the municipality of residence is observed. Table 2.1 sums up the main characteristics of these datasets.

2.3.1 Family insurance data

First, I use data provided by the CNAF, the Family branch of the French Social Security system, to get information on the supply of affordable collective childcare at the municipal level. Specifically, these data cover all EAJE facilities funded under the PSU scheme. For each municipality between 2007 and 2015, the data give the number of such facilities within each municipality and the number of childcare places they offer, as defined by their accreditation certificate, granted by the local authorities that specifies a maximum capacity for each facility.⁶

The Family branch of the French Social Security system also has data on the take-up of paid parental leave. Specifically, for each municipality from 2009 to 2018, this dataset gives the number of families who were entitled to either the CLCA or the PreParE in December of each year.⁷ In order to obtain to these allowances, families must submit an application and meet several criteria. This dataset therefore provides a relevant measure of the number of families that receive these parental leave allowances, as it only covers families who applied and are eligible.

⁶I exclude data on one département (Tarn), whose data would suggest that no EAJE-PSU facility existed in 2007, even though many municipalities had such facilities in 2008. In 2007, the Tarn département accounted for 0.6% of the total French population.

⁷Due to data issues related to a policy reform that took place in 2015, I restrict my analysis of this data to the 2009-2014 time-period: see Appendix 2.B.1 for additional details.

2.3.2 Labor market data

My labor market data are drawn from the *Déclarations Annuelles de Données Sociales* (DADS). By law,⁸ French employers have to fill in a DADS form for every employee subject to payroll taxes. The form contains detailed information about gross and net wages, days paid, hours paid, employer location (at municipality level), other job characteristics (beginning, duration and end of a period of employment and part-time employment), employer characteristics (industry, size, and region) and individual characteristics (age, gender and municipality of residence). In Appendix 2.B.2, I provide further details on how earnings and time worked are measured, and especially on how paid maternity leave is included in my measure of labor supply.

Specifically, I take advantage of two declination of these data. Firstly, I rely on the DADS panel, a longitudinal sample at rate 4.4% to track parents' labor supply and labor earnings from 2007 to 2015, thanks to an anonymized personal identifier based on their social security number that allows me to link this information to birth records. Secondly, I aggregate comprehensive cross-sectional DADS registers at the municipality level to recover earnings and hours paid to childminders and nannies from 2009 to 2015. Appendix 2.B.2 further details how I proceed and the limitations of these datasets.

2.3.3 Fertility data

My analysis also relies on birth records. Births are registered by an individual who was present at the time of birth, usually the father, but in some cases a doctor or a midwife. I again rely on two different versions of these records. Firstly, I take advantage of cross-sectional comprehensive birth records to compute the number of children born to women living in a given municipality in any year between 2005 and 2015, which gives an approximate measure of the trends in potential demand for childcare at a narrow geographical level. Secondly, I use on a longitudinal version of these records at the individual level extracted from the *Échantillon Démographique Permanent* (permanent demographic sample, EDP) to obtain information on the timing of births. Thanks to the NIR, this dataset can be merged with the longitudinal version of the DADS, and allows me to separate parents with young children from the rest of the population. Appendix 2.B.3 provides additional information as

⁸The absence of DADS as well as incorrect or missing answers are punished with fines.

to how I proceed and a few shortcomings of these data.

2.3.4 Data preparation

I first estimate the supply and potential demand for childcare at a narrow geographical level. For each municipality and each year the data provide information as to (i) the number of childcare places available in each municipality, and; (ii) the number of children born to mothers who live in the relevant geographical area. This allows me to compute a measure of the relative supply, i.e. the share of children with potential access to daycare center. Specifically, I define the relative childcare supply $S_{c,t}$, where c denotes municipality and t stands for a particular year, as the ratio:

$$S_{c,t} = \frac{N_{c,t}^{\text{places}}}{N_{c,t}^{\text{birth}} + N_{c,t-1}^{\text{birth}} + N_{c,t-2}^{\text{birth}}} \quad (2.1)$$

where $N_{c,t}^{\text{places}}$ is the number of EAJE-PSU childcare places in municipality c during year t , and $N_{c,t}^{\text{birth}}$ the number of children born to women who lived in c at time t . In other words, this measure assumes that children's place of residence does not change during the first three years of their life.

Figure 2.2 displays the trend in relative supply at the national level between 2007 and 2015. It increased by roughly 3.5 percentage points, and almost linearly over the period. An interesting feature of this continuous expansion of affordable childcare is its heterogeneity across geographical units. The map in Figure 2.3 displays the change in relative childcare supply level for each municipality from 2007 to 2015. It shows clearly that this moderate increase was concentrated in relatively few areas, where affordable childcare provision increased strongly, in contrast with most municipalities where the supply barely changed.

In a second step, I recover data at the individual level. I restrict the sample to individuals who experienced childbirth between 2005 and 2015, who therefore actually have children of the targeted age group at some point between 2007 and 2015, and to individuals between ages 20 and 60. As the municipality of residence is only observed in the labor market data, I further require that these individuals have been salaried employees at least once between 2002 and 2015.⁹

⁹Empirically, this is the case of 94% of parents throughout my time-period of interest. In Appendix 2.F.1, I show that under reasonable assumptions, this sample selection does not impede the identification of average treatment effects.

My analysis pays attention to the extensive margin of labor supply, which is crucial when considering mother's time allocation decisions. For individuals who are not found in my labor market data for a particular year,¹⁰ I impute zero labor earnings, and consider them to be outside the labor force.¹¹ As a result, I am able to decompose labor earnings responses between the extensive and intensive margins of labor supply on the one hand, and hourly wages on the other.

I finally merge this individual-level data with the geographical data on affordable childcare. This leaves me with 1.5 million observations of parents with children aged 0 to 2, covering 430 000 individuals. Table 2.2 gives summary statistics on the sample. The gender gap in labor outcomes is extremely salient: on average, while mothers of young children tend to be more educated, they earn only just over half the average earnings of their male counterparts. This gap is largely driven by labor supply decisions: among those in wage employment, the gender gap in hourly wages is much smaller yet still sizable, at about 15%.

2.4 Empirical analysis

2.4.1 Granular childcare expansions

My empirical approach leverages the granularity of national-level childcare expansion, i.e. the fact that (i) the smooth increase in childcare provision at the national level (Figure 2.2) is actually concentrated on a few municipalities where provision has increased sharply, in contrast with most municipalities where it has remained flat (Figure 2.3); and that (ii) among these municipalities in where childcare provision has increased massively, this rise is generally attributable to a single event, i.e. a sharp increase in affordable childcare provision between two consecutive years, for instance due to the opening of a new daycare center, rather than to a continuous increase over the years.

To make this granularity more salient, I first compute the maximum growth in relative childcare supply $S_{c,t}$ between two consecutive years in each municipality. Figure 2.4 displays the distribution of this maximal growth at the municipality level (weighted by the number of children aged 2 or less in each municipality as measured in 2007). In 2007, a quarter of children aged 2 or less lived in municipalities that

¹⁰I also drop observations with very low earnings or working time, see Appendix 2.B.4.

¹¹By contrast, all other observations correspond to individuals who are in employment.

experienced no growth in childcare supply of any kind between 2007 and 2015. In fact, these are mainly municipalities where the supply is actually nonexistent throughout the relevant time period, plus a handful of municipalities where the supply decreased due to the closure of a daycare facility. In municipalities that did experience growth, there is considerable heterogeneity in its maximum yearly magnitude: the 80th percentile of the distribution is 4 percentage points, the 90th percentile is 7.6 percentage points, but the 99th percentile is over 33 percentage points.

There is no obvious cut-off in the distribution. Nevertheless, I choose to partition municipalities into four treatment groups: those where the supply never increases (bottom 25%), those between the 25th and the 80th percentile, then those that rank between the 80th and the 90th percentile, and finally the top 10%. Dividing municipalities into separate groups according to the position in the distribution of a continuous variable is by no means straightforward; however this approach is somewhat similar to that of [Havnes and Mogstad \(2011a\)](#) who group municipalities according to their position below or above the median. Furthermore, and in contrast to theirs, my approach does not rely on heterogeneity across these groups. In [Appendix 2.C](#), I describe these groups in terms of pre-treatment observables, i.e. using data from the 2006 Census at the municipality level. The main lesson is that P90-P100 municipalities which are key to my identification strategy, are relatively small municipalities, with slightly less than 8,000 inhabitants in average.

I then define the timing of the childcare shock that corresponds to this maximum yearly growth. In municipalities that did experience positive growth, the definition is straightforward: the event takes place at the time when the relative childcare supply increases the most. For the bottom 25% of municipalities where the supply never increases, the counterfactual treatment time is drawn randomly in the distribution of actual treatment timings in the other groups.

[Figure 2.5](#) displays the average relative supply of affordable childcare over time within each treatment group, depending on the timing of the municipal childcare shock. In the never-treated group, this supply remains at around 0 from 2007 to 2015. For the three other groups, the figure clearly shows that within each group, the pre-shock level, the post-shock level and the size of the shock are very similar across municipalities with different timings of the shock itself. Basically, in the P25-P80 group of municipalities supply was 16-18% and increased by 1 percentage point; in the P80-P90 group supply was about 20%, and increased by 5 percentage points;

and in the P90-P100 group, pre-shock coverage was 15-20% and increased sharply 15 by percentage points. In this last group, this event corresponds typically to the opening of the first or the second facility in the municipality, which represents about 15 to 20 new childcare places.

2.4.2 Event-study analysis

I rely on differences in the timing of the childcare shock across municipalities that experience shocks of similar magnitudes to identify the causal impact of childcare expansions. Let y_{it} denote the annual earnings (resp. salaried employment dummy, working hours, hourly wages) of parent i at time t , living in municipality $c = c(i, t)$ that belongs to the treatment group $g = g(c)$.¹² In this within-group event-study setting, I estimate:

$$y_{it} = \alpha_{c(i,t)} + \sum_{g,\tau} \beta_{g\tau} \mathbb{1}\{t - E_{c(i,t)} = \tau, g(c(i,t)) = g\} + \sum_{g,T} \gamma_{gT} \mathbb{1}\{t = T, g(c(i,t)) = g\} + \epsilon_{it} \quad (2.2)$$

where α_c is a municipality-level fixed effect, E_c denotes the year of the childcare shock for municipality c , and ϵ_{it} is an idiosyncratic shock of mean 0. The $\beta_{g\tau}$ coefficients capture the dynamic effects of the childcare expansions and represent my parameter of interest. This parameter is identified thanks to variation in the timing of childcare expansions among municipalities that are affected by expansions of similar magnitude. In other words, it is not identified neither by (i) within-municipality differences in parental labor outcomes over time, as would be the case in an interrupted time-series approach, nor by (ii) differences in the evolution of parental outcomes between municipalities that are affected by an expansion and municipalities that do are not affected by an expansion, as would be the case in an usual difference-in-difference approach.

As noted by [Borusyak and Jaravel \(2017\)](#), Model 2.2 is underidentified. This is because (i) the inclusion of municipality fixed effects means that the time effects are only identified up to a constant; and more importantly, (ii) within each cohort defined by the timing of the treatment E_c , calendar time t and time-to-treatment $t -$

¹²Because the relevant municipality is the one in which parent i lives at time t , my approach takes into account families who may move from one municipality to another due to the opening of new childcare places.

E_c are colinear.¹³ This is actually a special case of the well-known underidentification problem of Age-Period-Cohort models, with age corresponding to time-to-treatment, period to calendar time, and cohort to the timing of the treatment. Due to this collinearity, the $\beta_{g\tau}$ coefficients are only identified up to a constant plus a linear trend.

To resolve this underidentification problem, [Borusyak and Jaravel \(2017\)](#) note that in settings where it is plausible to assume that (i) the treatment is exogenous conditional on unit (here: municipality) fixed-effects, and that (ii) there are no anticipation effects, coefficients belonging to the subset $(\beta_{g\tau})_{\tau < 0}$ should all be equal to 0. As a result, they suggest that Model 2.2 be estimated first, while setting two coefficients of the subset to 0, which is akin to APC modeling approach proposed by [Mason et al. \(1973\)](#). This makes it possible to test the hypothesis that other coefficients are also equal to 0. After ensuring that this no-pretrend assumption holds, they recommend estimating a semi-dynamic version of Model 2.2 in which all coefficients $(\beta_{g\tau})_{\tau < 0}$ are constrained to 0.

Lastly, they point out that when treatment effects are dynamic, i.e. when there is variation in the coefficients of the subset $(\beta_{g\tau})_{\tau \geq 0}$, the overall treatment effect is not identified by the canonical regression in which time-to-treatment dummies are replaced by a post-treatment dummy. This is because this regression weights long-run effects negatively: as a result, the estimator does not have the no-sign reversal property, so even the sign of the effect can be wrong. Instead, they recommend first fitting the semi-dynamic model, and then manually summing the coefficients of the subset $(\beta_{g\tau})_{\tau \leq 0}$, for instance with weights proportional to the sample size.

I follow their recommendation closely. My only departure is that as a first step, I do not normalize the pre-trend by setting two coefficients to 0. Instead, I apply a solution to the underidentification of APC models proposed by [Deaton and Paxson \(1994\)](#). Specifically, my approach basically involves imposing two normalizations on the pre-trend: (i) that on average, $\beta_{g\tau}$ coefficients before the event are equal to 0, i.e. $\sum_{\tau < 0} \beta_{g\tau} = 0$; and (ii) that the vector $(\beta_{g\tau})_{\tau < 0}$ is orthogonal to any linear time trend, i.e. $\sum_{\tau < 0} \tau \beta_{g\tau} = 0$.

Recent investigations of this approach show that these regressions can generate spurious results when treatment effects are heterogeneous across cohorts, as defined by the timing of the treatment ([Sun and Abraham, 2020](#)). In Appendix 2.F.3, I

¹³Municipality fixed effects can be replaced with cohort (time-of-the-treatment) fixed effects without changing the identification properties of the model.

show that moving to a correction based on a fully interacted model does not affect my results.

2.4.3 Instrumental variable approach

This event-study approach captures the consequences of childcare expansions without any reference to their magnitude. As a second step, I frame it into the fuzzy difference-in-difference approach developed by [Duflo \(2001\)](#) to rescale my estimates. In this setting, Model 2.2 is regarded as the reduced-form version of an instrumental variable regression, and is simply divided by the average magnitude of childcare expansions within the treatment group of the relevant municipality. Specifically, keeping the same notations, I estimate:

$$y_{it} = \kappa_{c(i,t)} + \lambda S_{c(i,t),t} + \sum_{g,T} \mu_{gT} \mathbb{1}\{t = T, g(c(i,t)) = g\} + \nu_{it} \quad (2.3)$$

while instrumenting the relative childcare supply S_{ct} by time-to-treatment interacted with treatment group dummies:

$$S_{ct} = \phi_c + \sum_{g,\tau \geq 0} \psi_{g\tau} \mathbb{1}\{t - E_c = \tau, g(c) = g\} + \sum_{g,T} \chi_{gT} \mathbb{1}\{t = T, g(c) = g\} + \omega_{ct} \quad (2.4)$$

The λ parameter can be interpreted at the individual level in an intention-to-treat sense: it corresponds to the effect on parents' labor outcomes of being offered a childcare place,¹⁴ for the restricted subset of parents who would not have been offered such a place before the local childcare expansion, but actually are due to the local childcare expansion. This interpretation rests on a Stable Unit Treatment Value Assumption which states that, within municipalities and conditional on whether they are assigned a childcare place or not, parents' labor supply decisions are independent of the assignment of childcare places to other families. In other words, there should be no peer effects in terms of labor supply, an assumption that is somewhat unrealistic ([Maurin and Moschion, 2009](#)). If this assumption fails, then my estimates should be interpreted as a more macro effect, incorporating social multipliers due to peer effects. In this case, when divided by 100, the λ parameter represents the causal effect of a one percentage-point increase in childcare provision at the municipality level, expressed as the fraction of children aged 2 or less covered by local EAJE-PSU

¹⁴regardless of whether they actually use it or not.

facilities, on parents' labor outcomes.

This fuzzy difference-in-difference framework has recently been investigated by econometricians who raise questions issues as to its ability to identify causal parameters of interest in realistic settings (de Chaisemartin and DHaultfœuille, 2018). In Appendix 2.F.3, I discuss these concerns and provide solutions to address them in the specific setting of this paper.

2.4.4 Identifying assumptions

My empirical framework is based on an event-study design. As such, it does not rely on differences between municipalities exposed to increases of different magnitudes in the supply of collective childcare. In other words, differences between the P90-P100 group and other treatment groups are not directly relevant for my approach: I do not assume that the assignment to any of these groups is exogenous.

Instead, key to my framework are differences in the timing of the shock across municipalities of the same group, and especially of the P90-P100 group. Specifically, my identifying assumption is that *within* the P90-P100 group, the counterfactual trend in parental labor earnings *absent* the local childcare shock is mean-independent of the year when this shock takes place.

The allocation of subsidies directed towards the opening of new childcare places may depart from this assumption if either (i) the decision of municipalities to apply to these subsidies, or (ii) the decision of local CAF offices to grant these subsidies are based on factors that also determine this counterfactual trend. As noted in Subsection 2.2.3, the attribution of these subsidies by local CAF offices is to a large extent only based on the level of the childcare coverage rate in the municipality (and not, for instance, its evolution). However, the municipalities decision to first apply and its determinants remain unknown.

To assess the plausibility of my identifying assumption in this context, I resort to Census data at the municipality level. This allows me to test whether, within treatment groups, the timing of the childcare shock is correlated with observable characteristics that could plausibly affect the counterfactual trend in parental labor supply. A substantial correlation would seriously question the validity of the mean-independence assumption upon which my framework rests.

I show that, within the P90-P100 treatment group, municipalities that are treated in the beginning of the 2007-2015 time-period are, in 2006, virtually in-

distinguishable from municipalities that are treated later on. Specifically, these municipalities differ very little in terms of labor market composition, couple and marriage formation and dissolution, and arrival of new residents. The only significant differences are that (i) municipalities with lower coverage rate in 2007 are treated earlier, which is consistent with Subsection 2.2.3; (ii) larger municipalities tend to be treated earlier; and (iii) municipalities with higher birth rates are treated later. Even so, these differences explain very little (2% at best) of the variance the timing of the local childcare shock. Appendix 2.F.2 details these findings.

My framework rests on another assumption, namely that the spatial level relevant to the childcare choices is the municipality level. Qualitative surveys indeed show that the municipality of residence is the only universal criterion for the allocation of childcare places (Onape, 2012). To the best of my knowledge, available quantitative data do not allow to scrutinize this claim. However, this assumption follows for instance the strategy used by Berger et al. (2021) who instrument individual childcare decisions by the exact same municipality-level collective childcare supply upon which I focus.

2.5 Parental earnings and labor supply effects

2.5.1 Graphical analysis

Figure 2.6 displays my estimates of the event-study approach to the labor earnings of mothers with children aged 0 to 2 respectively. First, it displays my estimates of the full dynamic model, in which the pretrend is normalized in line with the approach proposed by Deaton and Paxson (1994). Such estimates allow me to verify that all coefficients corresponding to time periods that predate the childcare expansions are not significantly different from 0, which is indeed the case. In other words, within each treatment group, and before they are treated, mothers' labor earnings evolve in parallel across municipalities with different timings of the childcare shock. This sustains the credibility of the no-pretrend assumption upon which my event-study approach is based.

This allows me to consider the estimates of the semi-dynamic model, i.e. the event-study model in which the pretrend is set to 0. I find that my estimates are never significantly different from 0 at the usual 95% level. My point estimates do not suggest that the effect becomes significantly positive over time, so these results

are not driven by short-run frictions.

An additional feature of my setting is that I can display estimates of the effect of non-existent or extremely small shocks to affordable collective childcare provision by considering the first two treatment groups. Consistent with the rationale, I find that such shocks have no effect on mothers' labor outcomes, which bears out the credibility of my identifying assumptions.¹⁵

Finally, I map these dynamic estimates into a single effect for each treatment group by summing the coefficients with weights proportional to the sample size. Table 2.3 displays my estimates, not only for labor earnings, but also for the potential margins of adjustment: labor force participation, working days, working hours per day and hourly wages. Consistent with my previous findings, I cannot detect any significant effect of the childcare shocks on mothers and fathers' labor earnings and labor supply. Moreover, these estimates are much more precise than my semi-dynamic estimates, so that economically significant effects can be largely ruled out: in the P90-P100 group, the aggregate effect of collective childcare expansions on mothers' salaried employment rate cannot exceed 2.6 percentage points.

2.5.2 Instrumental variable estimation

I then turn to the results of the related instrumental variable regression. These are merely the same results, but rescaled using the magnitude of the childcare shock as a first stage.

Table 2.4 displays my estimates. Consistent with my previous findings, I cannot detect any significant effect of affordable collective childcare provision on parents', and especially mothers' labor outcomes. While my standard errors may be quite large for overall labor earnings, they are sufficiently small for labor supply decisions at the extensive margin. Indeed, the upper bound of my 95% confidence intervals allows me to rule out effects larger than 5.3 percentage points, my point estimate being -1.7 percentage points. The same goes on for fathers, who are left virtually unaffected by the provision of affordable collective childcare. This result is however less surprising given that men's labor supply changes very little in respond to the arrival of children (Kleven et al., 2019).

¹⁵The negative effects in the never-treated group are not significant once the pre-trend is set to 0 (additional identification constraint in the event-study setting), and are not significant when aggregated in a single estimate. It is driven by a strong negative trend in the number of days worked.

To make sure that these results are driven by municipalities where collective childcare provision increased substantially, as opposed to others where childcare shocks are almost nonexistent, I restrict my sample to the P90-P100 group, and run the same regression. My results are in line with those from the whole sample: when only the P90-P100 treatment group is considered, the upper bound of the 95% confidence interval is 4.3 percentage points. This confirms that these results do indeed arise from the top of the distribution of childcare shocks.

Because a very large share of the overall growth in childcare coverage at the national level is driven by these childcare shocks, my estimates are to a large extent informative about the aggregate effect of the national plans. Appendix 2.D.1 shows that this implies that these plans had virtually no effect on the aggregate salaried employment rate of mothers.

Lastly, in Appendix 2.F.3, I discuss a variety of concerns about the validity of my identification strategy, e.g. the possible correlation with other policy changes, sampling issues or the validity of the fuzzy difference-in-difference setting. I show that these concerns do not affect my finding, i.e. that obtaining a childcare place does not lead to massive changes in mothers' labor outcomes.

2.6 Substitution across childcare solutions

I now investigate the crowding-out of other childcare solutions by the expansion of collective daycare provision, which might explain my null effects for maternal labor outcomes. To this end, I first consider the take-up of paid parental leave, and then investigate the demand for individualized childcare provided by childminders and nannies.

2.6.1 Paid parental leave

I use the CNAF dataset that provides information on the number of families receiving parental leave allowances at the municipality level as of 2009. Specifically, I divide this number by the number of children aged 2 or less to determine the share of parents receiving parental leave allowances for either a full-time or part-time parental leave.

I then apply my event-study analysis to these municipality-level data, on a restricted subset of municipalities that experienced a childcare shock between 2010

and 2014.¹⁶¹⁷ Figure 2.7 displays my estimates. Consistent with this rationale, I find that the expansion of affordable collective childcare facilities does not trigger any substantial change in the share of families with young children who receive parental leave allowances.

2.6.2 Individualized childcare

I rely on a cross-sectional and comprehensive version of the DADS dataset that provides information on earnings and hours paid to childminders and nannies, paid directly by households, as of 2009. Specifically, I aggregate hours at the municipality level for the entire 2009-2015 time period.

Childminders are subject to a strict regulation in terms of child-to-adult ratios, as are collective childcare facilities. Specifically, the law was changed in 2009, raising a childminders' maximum childcare capacity from 3 to 4 children.¹⁸ As a result, I propose a measure of the relative supply of formal individualized childcare at the municipality level as the total number of hours paid to childminders and nannies, multiplied by 4 and, divided by (i) the annual number of full-time employment spell hours (1820 hours), and (ii) the total number of children aged 0 to 2:

$$S_{c,t}^{\text{indiv}} = \frac{4H_{c,t}^{\text{indiv}}}{1820 \left(N_{c,t}^{\text{birth}} + N_{c,t-1}^{\text{birth}} + N_{c,t-2}^{\text{birth}} \right)} \quad (2.5)$$

where $H_{c,t}^{\text{indiv}}$ is the number of hours paid to childminders and nannies in municipality c during year t , and $N_{c,t}^{\text{birth}}$ is the number of children born to women who lived in c at time t . This measure approximates the concept of how many hours childminders and nannies work relative to how much they would be working if all children were under their care. It is not a perfect measure of this relative supply concept, however, because: (i) the legal four-children threshold includes the childminder's own children, who I cannot observe; and (ii) a childminders' maximum childcare capacity is fixed by an agreement quite similar to that of an EAJE facility, depends on their education, experience, and equipment (e.g. the number of rooms in their home). Four is the upper bound for this capacity. However, in 2014, the

¹⁶I weight the data by the number of children aged two or less as observed in 2007.

¹⁷Specifically, I implement the [Sun and Abraham \(2020\)](#) specification of the event-study design that allows for heterogeneous treatment effects across cohorts.

¹⁸Loi n° 2008-1330 du 17 décembre 2008 de financement de la sécurité sociale pour 2009

average number of children per childminder was 3.3 (Vroylandt, 2016) so that, while imperfect, this measure is not meaningless.

I then replicate my event-study analysis, with $S_{c,t}^{\text{indiv}}$ as the outcome, on a restricted subset of municipalities that experienced a childcare shock between 2010 and 2014.^{19,20} Figure 2.8 displays my estimates. I find that in the medium run, in municipalities that experienced the largest shocks on collective childcare supply, substitution effects dominate: demand for childminders and nannies drops substantially. The magnitude of my estimates, about 13 percentage points, is almost equal to the magnitude of the corresponding collective childcare expansions (14 p.p.). This suggests sizeable crowding-out effects are at play: in other words, childcare expansions tend to shift families away from costly individualized childcare solutions.

On top of explaining the null effect of collective childcare on maternal labor supply, these substitution effects are crucial for the evaluation of the policy at stake. Indeed, in the French context in which all formal childcare solutions are subsidized through different channels, taking them into account changes the cost of a collective childcare place for public finances quite dramatically. Appendix 2.D.2 develops this point and shows that, when the crowding-out of individualized childcare solutions is taken into account, the estimated cost of the policy is divided by 2.7.

Lastly, in Appendix 2.F.4, I assess the robustness of these results to various concerns regarding the validity of my identification strategy, e.g. correlation with other policy changes or division bias. I find my results on parental leave take-up and demand for individualized childcare to be unaffected by these issues.

2.7 Conclusion

In this paper, I leverage differences across French municipalities in the timing of collective childcare expansions to identify the causal impact of affordable collective childcare on parents' labor outcomes. Applying an event-study framework to a combination of administrative records, I show that such expansions did not trigger any substantial change in the labor earnings and labor supply of parents with children in the targeted age groups. Interpreted as a local average treatment effect (LATE), my instrumental variable estimates suggest that, among mothers who obtained a collec-

¹⁹I weight the data by the number of children aged 2 or less as observed in 2007.

²⁰Here again, I implement the Sun and Abraham (2020) specification of the event-study design that allows for heterogeneous treatment effects across cohorts.

tive childcare place thanks to these expansions, this treatment did not strengthen labor market attachment. This is because the expansion of affordable collective childcare did not make mothers any less likely to benefit from paid parental leave. I provide evidence that instead, these expansions resulted first and foremost in a substantial crowding-out of individualized and more costly childcare solutions.

As these estimates are only informative about the choices of parents who were offered a childcare place under the national plans that I investigate, these results do not contradict the intuition that the lack of affordable childcare solutions may prevent some mothers from entering the workforce when they have young children (see Appendix 2.A.4). Instead, they draw attention to the selection of recipients of these newly created childcare places, who, my results suggest, would have otherwise relied on other formal childcare solutions.

Two mechanisms may explain these results. The first one deals is into application: it might be that families who would benefit most from a place are less likely to apply, due possibly to heterogeneity in preferences, exposure to social norms or heterogeneous returns on time spent in the labor market. For instance, strong cultural norms regarding childcare provided by mothers may prevent some families from applying for a collective childcare place, even though obtaining a place would actually change their work-family balance. The second is selection into treatment: in this setting, among actual applicants, childcare place may be offered preferentially to families who will benefit less from them. Survey data suggests, for instance, that because one of their roles is to foster a better work-family balance, about two thirds of EAJE-PSU facilities give higher priority to families in which both parents hold a full-time job (Onape, 2012). Conditioning treatment on actual observed outcomes, instead of unobserved treatment effects would then result in inefficiencies (Yamaguchi et al., 2018). Disentangling the two mechanisms therefore has relevant policy implications, but requires additional data on childcare preferences, application and selection into collective facilities.

Lastly, this empirical policy evaluation exercise does not consider how childcare choices affect children themselves, whose long-term outcomes may be substantially affected. Indeed, early childcare choices may affect children's health and early learning, thereby affecting their future socialization, education and labor market prospects (see e.g. Havnes and Mogstad, 2011b; Berger et al., 2021). These potential lifecycle benefits must be taken into account to achieve a meaningful normative analysis of these policies.

References

- Allègre, G., Simonnet, V., and Sofer, C. (2015). Child care and labour market participation in france: Do monetary incentives matter? *Annals of Economics and Statistics*, (117/118):115–139.
- Andresen, M. E. and Havnes, T. (2019). Child care, parental labor supply and tax revenue. *Labour Economics*, 61:101762.
- Asai, Y., Kambayashi, R., and Yamaguchi, S. (2015). Childcare availability, household structure, and maternal employment. *Journal of the Japanese and International Economies*, 38:172 – 192.
- Baker, M., Gruber, J., and Milligan, K. (2008). Universal child care, maternal labor supply, and family wellbeing. *Journal of Political Economy*, 116(4):709–745.
- Bassok, D., Fitzpatrick, M., and Loeb, S. (2014). Does state preschool crowd-out private provision? the impact of universal preschool on the childcare sector in oklahoma and georgia. *Journal of Urban Economics*, 83:18 – 33.
- Bauernschuster, S. and Schlotter, M. (2015). Public child care and mothers' labor supply Evidence from two quasi-experiments. *Journal of Public Economics*, 123(C):1–16.
- Berger, L. M., Panico, L., and Solaz, A. (2021). The Impact of Center-Based Child-care Attendance on Early Child Development: Evidence From the French Elfe Cohort. *Demography*, 58(2):419–450.
- Berlinski, S. and Galiani, S. (2007). The effect of a large expansion of pre-primary school facilities on preschool attendance and maternal employment. *Labour Economics*, 14(3):665–680.
- Bick, A. (2016). The Quantitative Role of Child Care for Female Labor Force Participation and Fertility. *Journal of the European Economic Association*, 14(3):639–668.
- Blau, D. and Currie, J. (2006). Chapter 20 pre-school, day care, and after-school care: Who's minding the kids? volume 2 of *Handbook of the Economics of Education*, pages 1163 – 1278. Elsevier.
- Borusyak, K. and Jaravel, X. (2017). Revisiting event study designs. mimeo.
- Carta, F. and Rizzica, L. (2018). Early kindergarten, maternal labor supply and children's outcomes: Evidence from italy. *Journal of Public Economics*, 158:79 – 102.
- Cascio, E. U. (2009). Maternal labor supply and the introduction of kindergartens into american public schools. *Journal of Human Resources*, 44(1):140–170.
- Choné, P., Le Blanc, D., and Robert-Bobée, I. (2004). Offre de travail féminine et garde des jeunes enfants. *Économie et Prévision*, 162(1):23–50.
- Clément, J. and Aho, F. (2018). Atlas des eaje – exercice 2015. CNAF.
- Connelly, R. (1992). Self-employment and providing child care. *Demography*, 29(1):17–29.
- de Chaisemartin, C. and DHaultfoeuille, X. (2018). Fuzzy Differences-in-Differences. *The Review of Economic Studies*, 85(2):999–1028.

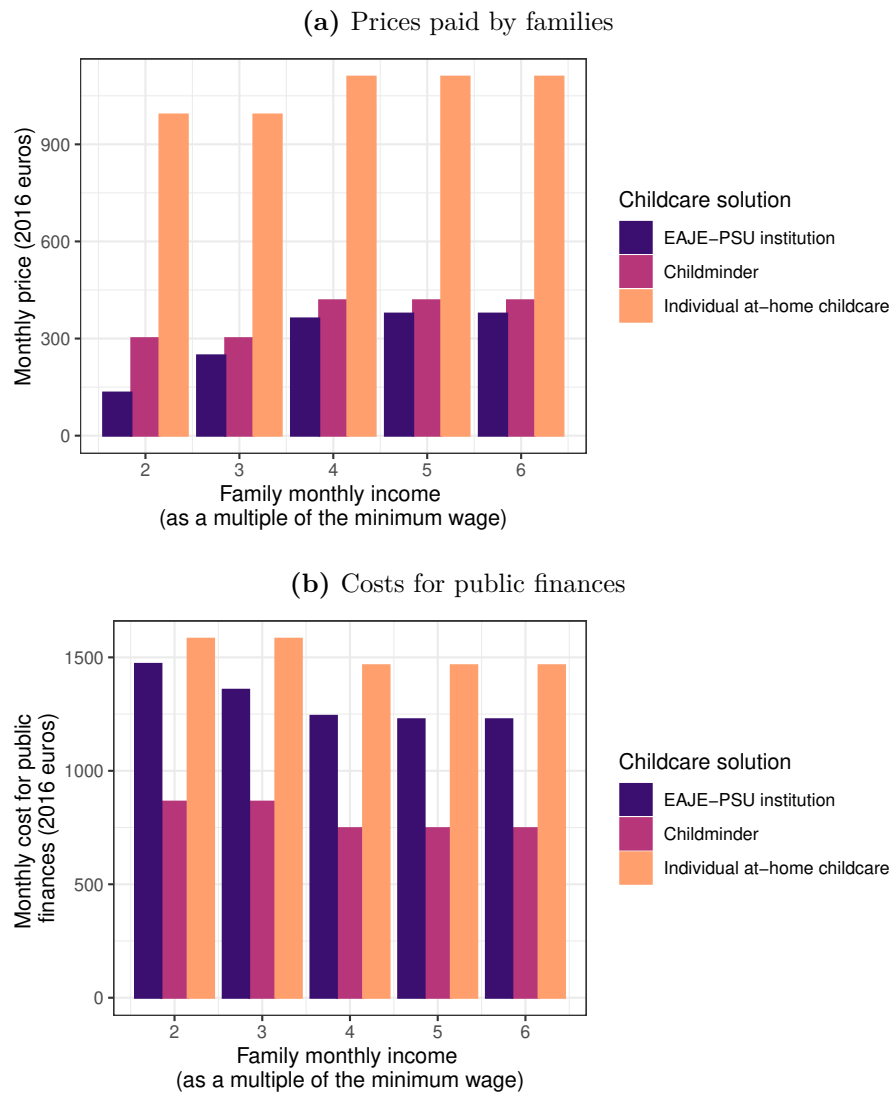
- Deaton, A. S. and Paxson, C. H. (1994). Saving, growth, and aging in taiwan. In Wise, D. A., editor, *Studies in the Economics of Aging*, pages 331–362. University of Chicago Press.
- Domeij, D. (2013). Should Day Care be Subsidized? *Review of Economic Studies*, 80(2):568–595.
- Duffo, E. (2001). Schooling and labor market consequences of school construction in indonesia: Evidence from an unusual policy experiment. *American Economic Review*, 91(4):795–813.
- Fitzpatrick, M. D. (2010). Preschoolers enrolled and mothers at work? the effects of universal prekindergarten. *Journal of Labor Economics*, 28(1):51–85.
- Galtier, B. (2011). Larbitrage entre emploi et inactivité des mères de jeunes enfants : le poids des contraintes familiales, professionnelles et sociétales sur les modes d'accueil des enfants. *Économie et statistique*, 477.
- Gathmann, C. and Sass, B. (2018). Taxing Childcare: Effects on Childcare Choices, Family Labor Supply, and Children. *Journal of Labor Economics*, 36(3):665–709.
- Gelbach, J. B. (2002). Public schooling for young children and maternal labor supply. *American Economic Review*, 92(1):307–322.
- Givord, P. and Marbot, C. (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.
- Goux, D. and Maurin, E. (2010). Public school availability for two-year olds and mothers' labour supply. *Labour Economics*, 17(6):951–962.
- Havnes, T. and Mogstad, M. (2011a). Money for nothing? universal child care and maternal employment. *Journal of Public Economics*, 95(11):1455–1465.
- Havnes, T. and Mogstad, M. (2011b). No child left behind: Subsidized child care and children's long-run outcomes. *American Economic Journal: Economic Policy*, 3(2):97–129.
- Heckman, J. J. (1974). Effects of child-care programs on women's work effort. *Journal of Political Economy*, 82(2):S136–S163.
- Hotz, V. J. and Xiao, M. (2011). The impact of regulations on the supply and quality of care in child care markets. *American Economic Review*, 101(5):1775–1805.
- IGAS/IGF (2017). La politique d'accueil du jeune enfant revue de dépenses. Rapport IGAS N°2017-019R / IGF N°2017-M-014, Inspection Générale des Affaires Sociales.
- Joseph, O., Pailhé, A., Recotillet, I., and Solaz, A. (2013). The economic impact of taking short parental leave: Evaluation of a french reform. *Labour Economics*, 25(C):63–75.
- Kleven, H., Landais, C., Posch, J., Steinhauer, A., and Zweimüller, J. (2020). Do family policies reduce gender inequality? evidence from 60 years of policy experimentation. Working Paper 28082, National Bureau of Economic Research.
- Kleven, H., Landais, C., and Sogaard, J. E. (2019). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.

- Lequien, L. (2012). The impact of parental leave duration on later wages. *Annals of Economics and Statistics*, (107/108):267–285.
- Mason, K. O., Mason, W. M., Winsborough, H. H., and Poole, W. (1973). Some methodological issues in cohort analysis of archival data. *American Sociological Review*, 38(2):85–97.
- Maurin, E. and Moschion, J. (2009). The Social Multiplier and Labor Market Participation of Mothers. *American Economic Journal: Applied Economics*, 1(1):251–272.
- Maurin, E. and Roy, D. (2008). L’effet de l’obtention d’une place en crèche sur le retour à l’emploi des mères et leur perception du développement de leurs enfants. Docweb 0807, CEPREMAP.
- Michalopoulos, C., Robins, P. K., and Garfinkel, I. (1992). A Structural Model of Labor Supply and Child Care Demand. *Journal of Human Resources*, 27(1):166–203.
- Morrissey, T. W. (2017). Child care and parent labor force participation: a review of the research literature. *Review of Economics of the Household*, 15(1):1–24.
- Müller, K.-U. and Wrohlich, K. (2020). Does subsidized care for toddlers increase maternal labor supply? evidence from a large-scale expansion of early childcare. *Labour Economics*, 62:101776.
- Nix, E. and Andresen, M. E. (2019). What Causes the Child Penalty? Evidence from Same Sex Couples and Policy Reforms. Discussion Papers 902, Statistics Norway, Research Department.
- Nollenberger, N. and Rodríguez-Planas, N. (2015). Full-time universal childcare in a context of low maternal employment: Quasi-experimental evidence from Spain. *Labour Economics*, 36(C):124–136.
- OECD (2011). Doing better for families.
- OECD (2016). Family indicators. OECD Social and Welfare Statistics (database).
- Onape (2010). L’accueil du jeune enfant en 2009. Rapport, Observatoire national de la petite enfance.
- Onape (2012). L’accueil du jeune enfant en 2011. Rapport, Observatoire national de la petite enfance.
- Onape (2017). L’accueil du jeune enfant en 2016. Rapport, Observatoire national de la petite enfance.
- 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*, (156):79–109.
- Pora, P. and Wilner, L. (2019). Child penalties and financial incentives: Exploiting variation along the wage distribution. Documents de Travail de l’Insee - INSEE Working Papers G2019/08, Insee.
- Sun, L. and Abraham, S. (2020). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*.
- Villaume, S. and Legendre, E. (2014). Modes de garde et d’accueil des jeunes enfants en 2013. Études et Résultats 896, Drees.

- Vroylandt, T. (2016). Les assistantes maternelles gardent 8% denfants de plus en 2014 quen 2010. *Études et Résultats* 978, Drees.
- Wilner, L. (2016). Worker-firm matching and the parenthood pay gap: Evidence from linked employer-employee data. *Journal of Population Economics*, 29(4):991–1023.
- Yamaguchi, S., Asai, Y., and Kambayashi, R. (2018). Effects of subsidized childcare on mothers labor supply under a rationing mechanism. *Labour Economics*, 55:1 – 17.

Figures


Figure 2.1 – Childcare prices along the income distribution



Monthly price paid by families and monthly cost for public finances along the income distribution, by choice of childcare solution.

Source. CNAF, case-study estimates (Onape, 2017).

Figure 2.2 – Relative supply EAJE-PSU affordable collective childcare at the national level from 2007 to 2015

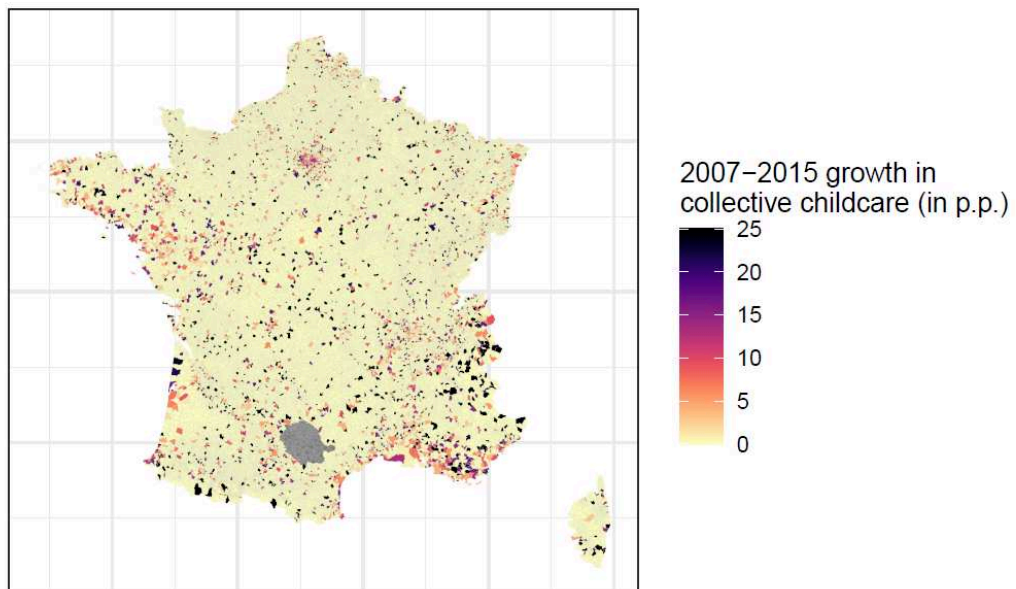


Estimates of the ratio EAJE-PSU childcare places offered to children aged 2 or less in metropolitan France (mainland France and Corsica).

Note. Data on the Tarn département are omitted.

Source. EAJE-PSU records, CNAF. Birth records, Insee.

Figure 2.3 – Spatial distribution of the 2007-2015 growth in relative supply of EAJE-PSU affordable collective childcare

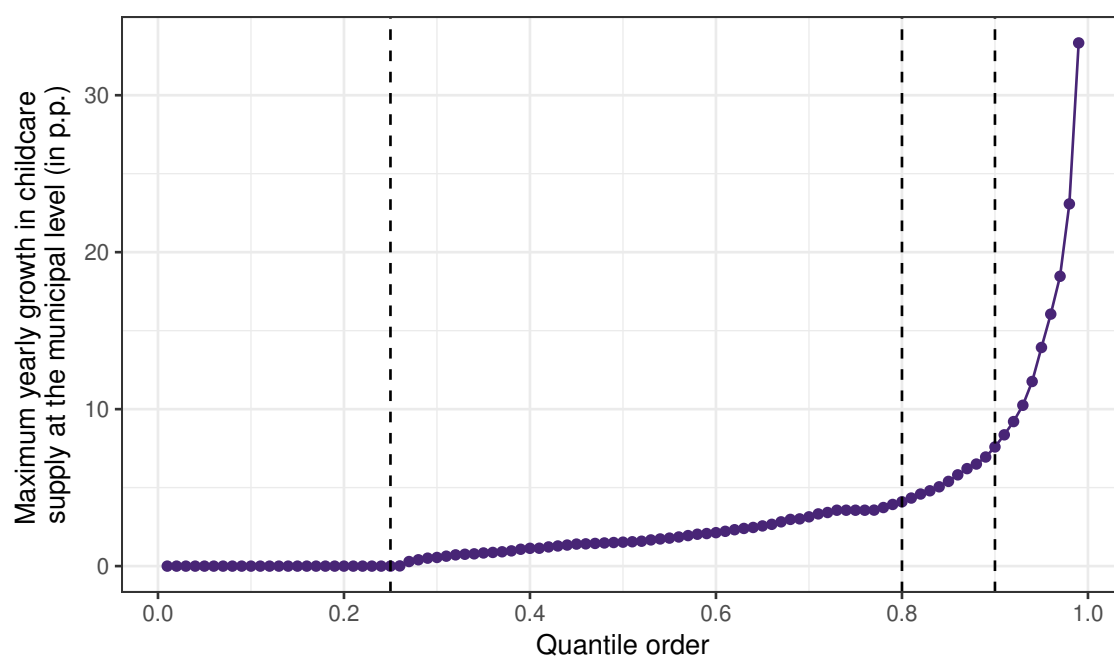


Estimates of the 2007-2015 growth in the ratio of EAJE-PSU places offered to children aged 2 or less at municipality level.

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records, CNAF. Birth records, Insee.

Figure 2.4 – Distribution of maximum annual within-municipality growth in affordable collective childcare coverage

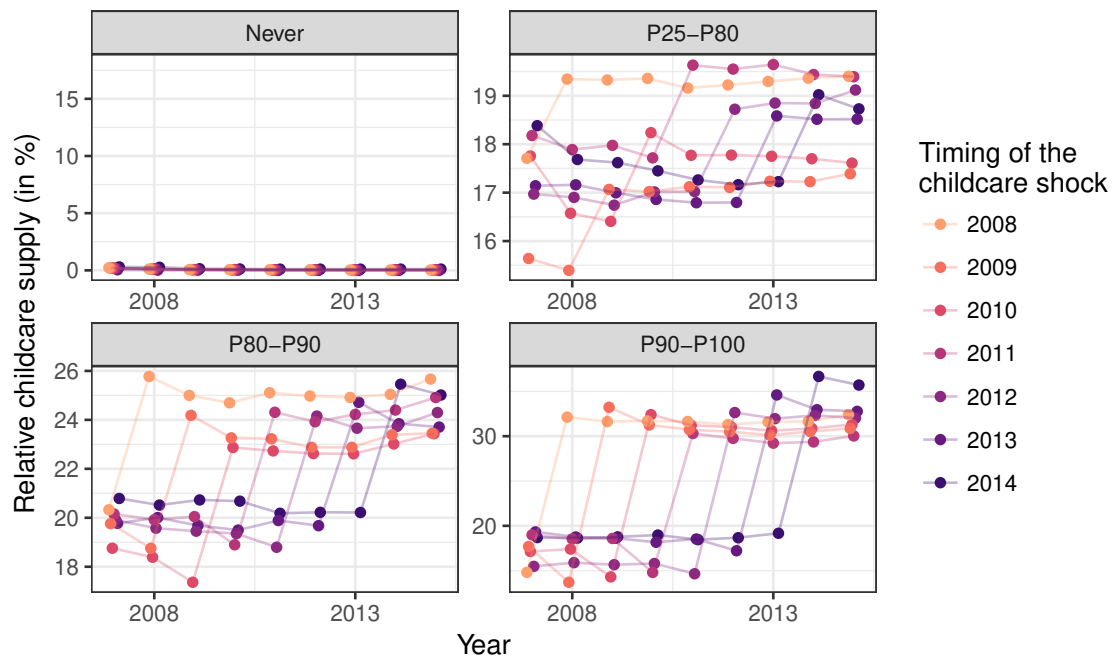


Estimates of the highest annual within-municipality growth in the ratio EAJE-PSU childcare places offered to children aged 2 or less.

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records, CNAF. Birth records, Insee.

Figure 2.5 – Relative supply of EAJE-PSU affordable childcare, by treatment group and timing of the childcare shock




Estimates of the ratio of EAJE-PSU childcare places offered to children aged 2 or less at the municipality level.

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records, CNAF. Birth records, Insee.

Figure 2.6 – Event-study estimates of the impact of the childcare shock on mothers' labor earnings, by treatment group




Event-study estimates of the effect of childcare shocks on mothers' labor earnings (Model 2.2).

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records, CNAF. Birth records and DADS-EDP panel, Insee.

Figure 2.7 – Event-study estimates of the impact of the childcare shock on paid parental leave take-up, by treatment group




Event-study estimates of the effect of childcare shocks on the share of families receiving parental leave allowances.

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU and PAJE records, CNAF. Birth records, Insee

Figure 2.8 – Event-study estimates of the impact of the childcare shock on the supply of individualized childcare, by treatment group



Event-study estimates of the effect of childcare shocks on individualized childcare by childminders and nannies.

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records, CNAF. Birth records and comprehensive DADS records, Insee

Tables

Table 2.1 – Data description

Dataset	Source	Main variables	Individual identifier	Municipality identifier
EAJE records	CNAF	# childcare places		✓
PAJE records	CNAF	# families receiving parental leave benefits		✓
DADS panel	Insee	Earnings, days and hours worked	✓	✓
DADS comprehensive records	Insee	Earnings, days and hours worked, detailed occupation		✓
Birth records	Insee	Date of birth		✓
EDP panel	Insee	Date of birth of parents' children, education	✓	

Table 2.2 – Summary statistics

	Mothers		Fathers	
	Mean	Standard Deviation	Mean	Standard Deviation
# Observations	740,412		775,658	
# Individuals	212,108		221,335	
<i>a. Individual characteristics</i>				
Age	31.4	5.1	34.1	6.2
Number of children [*]	1.8	0.9	1.8	1.0
Higher education ^{**}	0.22	0.17	0.18	0.15
Lower secondary education ^{**}	0.09	0.08	0.14	0.12
<i>b. Treatment rate</i>				
Childcare supply	15.0	14.8	15.0	14.5
<i>c. Labor outcomes</i>				
Earnings (2015€)	10,760	12,490	19,460	19,600
Employment	0.67	0.47	0.82	0.38
Days worked	317	125	345	122
Hours per day	4.0	1.3	4.8	1.1
Hourly wages (2015€)	12.1	5.8	14.1	9.1

^{*} Among individuals born in October. ^{**} Among those with available information. *Note.* Data regarding the Tarn département are omitted. *Source.* EAJE-PSU records, CNAF. Birth records and DADS-EDP panel, Insee.

Table 2.3 – Event-study estimates of the impact of childcare expansions on parents' labor outcomes

Treatment group	Childcare supply (p.p.)	Labor earnings (2015 euros)	Employment (p.p.)	Days	Hours per day	Hourly wages (2015 euros)
<i>Mothers with children aged 0 to 2</i>						
Never	0.04 (0.05)	-287.58 (154.23)	-0.49 (0.7)	-4.69 (2.12)	-0.014 (0.023)	-0.024 (0.073)
P25-P80	1.98 (0.27)	51.51 (139.63)	0.6 (0.6)	2.3 (1.86)	0.016 (0.019)	-0.031 (0.076)
P80-P90	5.03 (0.38)	128.88 (337.47)	0.97 (1.12)	0.43 (3.57)	-0.023 (0.041)	-0.224 (0.171)
P90-P100	17.55 (0.78)	348.57 (301.11)	0.46 (1.1)	1.02 (3.52)	0.002 (0.037)	0.266 (0.156)
<i>Fathers with children aged 0 to 2</i>						
Never	0 (0.05)	268.26 (215.56)	0.01 (0.51)	2.41 (1.58)	-0.006 (0.016)	0.065 (0.102)
P25-P80	1.97 (0.27)	257.38 (196.45)	0.03 (0.51)	0.96 (1.5)	0.028 (0.015)	0.08 (0.09)
P80-P90	5.13 (0.38)	-840.95 (495.65)	1.22 (0.94)	-1.79 (3.34)	0.004 (0.029)	-0.826 (0.243)
P90-P100	16.94 (0.7)	-226.41 (443.8)	1.23 (0.89)	-4.15 (2.92)	-0.017 (0.026)	-0.07 (0.242)

Dependent variable. EAJE-PSU childcare supply and parents' labor outcomes. *Explanatory variables.* Time-to-event and calendar-time dummies, interacted with treatment group, plus municipality fixed effects. Standard errors are clustered at the municipality level. *Note.* Data regarding the Tarn département are omitted. *Source.* EAJE-PSU records, CNAF. Birth records and DADS-EDP panel, Insee.

Table 2.4 – Instrumental variable estimates of the impact of affordable collective childcare on parents’ labor outcomes, by gender

Age of youngest child	Labor earnings (2015 euros)	Employment (p.p.)	Days	Hours per day	Hourly wages (2015 euros)
<i>Mothers</i>					
<i>a. Full sample</i>					
0–2	407.48 (894.34)	-1.73 (3.58)	-2.66 (11.23)	-0.066 (0.126)	0.681 (0.473)
<i>b. P90-P100 treatment group</i>					
0–2	174.44 (944.92)	-3.07 (3.78)	-9.41 (11.82)	-0.027 (0.134)	0.722 (0.486)
<i>Fathers</i>					
<i>a. Full sample</i>					
0–2	603.78 (1395.13)	3.02 (2.8)	-2.64 (9.4)	0.005 (0.088)	0.053 (0.671)
<i>b. P90-P100 treatment group</i>					
0–2	-91.51 (1461.17)	2.72 (2.9)	-1.85 (9.84)	-0.018 (0.093)	-0.09 (0.715)

Dependent variable. Parents’ labor outcomes. *Explanatory variables.* Childcare supply and calendar-time dummies interacted with treatment group, plus municipality fixed effects. Childcare supply is instrumented by time-to-event dummies interacted with treatment group. Standard errors are clustered at the municipality level. *Note.* Data regarding the Tarn département are omitted. *Source.* EAJE-PSU records, CNAF. Birth records and DADS-EDP panel, Insee.

2.A Institutional background

2.A.1 EAJE-PSU institutions

EAJE facilities are strictly regulated, in accordance with the Public Health Code, and cannot operate without authorization from local authorities (either at the département level for private facilities, or at municipal level for public ones), after an accreditation by the Maternal and Child Health Services. For each facility, this authorization defines a maximum capacity in terms of the number of childcare places.

The PSU funding covers 66% of the hourly cost of childcare, after families' contributions have been deducted. To obtain it, an EAJE facility must meet several requirements: (i) it has been authorized to open by the relevant authorities; (ii) its daycare places are open to all families; (iii) its pricing is based on a national fee schedule that makes this type of childcare particularly affordable to families; and (iv) it has signed an agreement on targets and management practices with the local CAF office.

In the national fee schedule, the upper bound of the *hourly* price paid by families is about 0.06% of their total *monthly* income, with a lower and an upper threshold on the total fees. A general rule of thumb is that the direct cost for parents of a full-time childcare place is between 5% and 10% of household income (IGAS/IGF, 2017). In 2015, the average hourly price that families paid was €1.80 (Clément and Aho, 2018). By contrast, other formal childcare solutions, i.e. childminders or at-home childcare were much more expensive, especially for families at the lowest end of the income distribution.²¹

To achieve these low prices, EAJE-PSU childcare facilities are heavily subsidized. In 2015, the average total hourly cost of an EAJE place was €8.86 (Clément and Aho, 2018), the average operating cost of a full-time EAJE place was €15 000, of which the Family branch of the Social security contributed up to 44.4%, local authorities 19.1%, other public stake-holders 18% and families 18.1% (Onape, 2017).²² Overall, the operating costs of EAJE-PSU facilities in 2015 amounted to €6 billion.

Lastly, EAJE facilities were opened for 222 days a year on average in 2015, which amounts to around 5 days a week excluding July and August (Onape, 2017),

²¹A reason for this is that individualized childcare solutions are subsidized through tax credits and tax rebates that make them particularly appealing to families that pay large income taxes.

²²Childcare subsidies are not restricted to EAJE-PSU facilities, as other forms of childcare are also heavily subsidized.

and for 11 hours a day on average. With respect to an ideal situation of full-time occupancy of all childcare places throughout the year, the occupancy rate was about 70%. These quantities varied very little throughout my period of interest.

2.A.2 National plans

To benefit from subsidies related to the national plans, the instigators of a project, usually a municipality, have to apply to the relevant local CAF office. National guidelines and criteria edicted by the CNAF are then used to assess the eligibility of each project. Unfortunately, the detail of these guidelines and criteria and their changes over time is not available. In 2013, the national guidelines stated that local CAF offices only rank projects according to the coverage rate, i.e. the number of formal childcare places relative to the number of children aged 3 or less.²³ As a result, municipalities with low coverage rate are given a higher priority. Additional criteria can be used to offer additional subsidies to applicants, e.g. if municipalities have a particularly low coverage rate or are relatively poor so that local taxes are less likely to cover the costs of the project.

These investment plans represent a substantial burden for public finances: in 2016, it was estimated that since the launch of the first plan in 2000, spending on the *plans crèches* had totaled €2 billion, not counting the annual operating costs of the childcare places created (IGAS/IGF, 2017). The average cost of creating a childcare place was estimated about €27 000 in 2009, 29% of which was financed through the national expansion plans (Onape, 2010).

²³Lettre circulaire n° 2013-149 de la Direction des politiques familiales et sociale

2.A.3 Parental leave policies

Various parental benefits were merged in 2004 into the *Prestation d'Accueil du Jeune Enfant* (PAJE). The PAJE comprises means-tested lump-sum payment after a birth (*prime de naissance*), monthly means-tested benefits (*allocations familiales*), a childcare subsidy (*Complément libre choix du Mode de Garde* (CMG)), and some benefits that may be granted when a parents interrupt his or her career or opts to work part-time (previously *Complément Libre Choix d'Activité* (CLCA) and now *Prestation Partagée d'Éducation de l'enfant* (PreParE)). Additionally, parents are entitled by law to extend the duration of their parental leave if they are not offered a formal childcare place.²⁴

These child benefits date back to 1985 and were introduced with the creation of *Allocation Parentale d'Éducation* (APE) initially restricted to mothers of 3 or more children. The APE was extended to mothers of 2 children in 1994, and was replaced by the CLCA in 2004, becoming effective with the first birth and providing a fixed non-means-tested monthly amount for the maximum duration of 6 months. The CLCA was replaced in 2015 by PreParE that introduced incentives to split parental leave between both parents. Contrary to Sweden, for instance, the benefits do not depend on parents' past income: they amount to approximately €400 per month in the case of career interruption and to nearly €200 in the case of 80% part-time work. Several papers have shown that these benefits encourage some mothers to reduce their labor supply (Choné et al., 2004; Piketty, 2005; Lequien, 2012; Joseph et al., 2013).

²⁴Article L531-4 of the French Social Security Code.

2.A.4 Childcare preferences and choices


On top of these institutional pushes towards increasing provision of affordable, formal collective childcare, there is strong demand from parents for these services. In 2010, Insee implemented a complementary module to the French Labor Force Survey (*Enquête Emploi*) devoted to the work-family balance. In this survey, 1999 individuals with children under age 3 were asked what type of childcare solution they thought was ideal for children the same age as their youngest child; what was their actual choice of childcare arrangement; what kinds of constraints they met when making this choice; and lastly, whether this choice impacted their labor supply decisions. I take advantage of these data to shed further light on potential demand for the kind of childcare solutions examined in this paper.

Firstly, daycare centers are among parents' preferred childcare arrangements: over 25% of both mothers and fathers of children under age 3 mention it as the ideal childcare solution (Figure 2.A.1). While a slightly higher proportion of parents view childcare by parents as ideal (about 30%), no other childcare solution is more frequently mentioned by parents as their preferred option.

Secondly, while 67% of parents who indeed rely on this childcare arrangement view it as ideal, only 32% of parents who mention it as their preferred solution actually use it (Figure 2.A.2). Additionally, these are most likely to report difficulties in accessing childcare: 31% of them report difficulties of this kind (Figure 2.A.3), among whom 70% mention the lack of availability of their desired childcare solution as the main problem encountered.

Lastly, among mothers of young children who are not working full-time and who do not use daycare centers, those who consider them to be the ideal childcare solution are the most likely to report that either insufficient availability, or the cost of childcare impacted their labor supply decision (Figure 2.A.4).


Figure 2.A.1 – Ideal childcare solution reporting by parents of children under age 3



Ideal childcare solution for children of the same as their youngest child, reported by parents of children under age 3.

Source. LFS complimentary module 2010, Insee.


Figure 2.A.2 – Actual childcare choices of parents of children under age 3



Actual childcare solution used by parents of children under age 3, by preferred childcare solution.

Source. LFS complimentary module 2010, Insee.


Figure 2.A.3 – Difficulties in childcare access as reported by parents of children under age 3



Share of parents of children under age 3 who report difficulties in access to childcare, by preferred childcare solution.

Source. LFS complimentary module 2010, Insee.

Figure 2.A.4 – Impact of childcare availability on labor supply decisions, as reported by parents of children under age 3



Share of parents of children under age 3 who do not work full-time who report reducing their labor supply due to a lack of appropriate childcare, by gender and preferred childcare solution.

Source. LFS complimentary module 2010, Insee.

2.B Data

2.B.1 Family insurance data

The parental leave dataset provided by CNAF suffers from a few issues. First, it does not allow to distinguish between full-time and part-time parental leave-takers. Second, it does not provide information on which parent took the leave. Third, in 2015 there were no data on the take-up of PreParE – the newly created parental leave allowance, available to parents of children born in 2015 or later. For this reason, when investigating the take-up of parental leave benefits, I rely solely on data that cover the 2009-2014 period, and only consider the take-up of the CLCA allowance. Lastly, for reasons of personal data protection, the dataset does not provide information on municipalities where fewer than 5 families took the leave. To the extent that childcare facilities are generally located in municipalities where number are well above this threshold, this does not appear to be a major obstacle to my analysis.

2.B.2 Labor market data

Earnings and working time measures

Earnings My measure of labor earnings is based on net annual earnings. This measure aggregates all wages paid to an individual, including performance pay and bonuses, annual leave, in-kind benefits, the share of severance payments that exceeds the legal minimum, and early retirement benefits (to the extent that these benefits exceed an amount approximately equal to the minimum wage) but excludes stock-options. Social security and public pension contributions, unemployment contributions and other contributions, including two flat-rate taxes on earned income (CSG and CRDS), are subtracted from this amount to compute our measure of net annual earnings. I thus measure earnings before income tax but after some transfers.

Maternity leave allowances are paid by the Social Security administration, and as such are not part of my measure of earnings. They may, however, be paid through the employer (*subrogation*), in which case the employer pays the employee the equivalent of maternity leave allowances during her maternity leave, and is later reimbursed by the Social Security administration. The maternity leave allowances that the employer advanced are subtracted from my measure of earnings. As the amount is reimbursed after the maternity leave has begun, the woman's decline in earnings may occur a few weeks after the start of her maternity leave. Because I consider annual earnings, this problem is restricted to births that occur at the end of the calendar year.

Lastly, in some firms the employer may be required under a collective agreement to complement earnings during maternity or sick leaves on top of the allowances paid by the Social Security. This complement is counted by the DADS as labor earnings.

Days In the DADS dataset, days paid refers to the duration of an employee's presence in a firm's workforce within a given year. As a result, maternity leave, sick leave, or paid annual leave are part of this measure of days, whereas a period of unemployment between two distinct employment spells is not. Additionally, this measure of days is capped at 360.

Hours In the DADS dataset, hours paid refers to hours for which the worker is paid under their labor contract. The data on hours is reported by employers when

they fill out payroll tax forms. Before making the data available, Insee performs three checks:

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

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

For workers whose pay does not depend on the time worked, but who do not belong to any of the above-mentioned occupations, i.e., typically highly-qualified personnel working in a "day rate" (*"forfait-jour"*), employers provide the number of days only. A number of hours is ascribed to these observations based on the legal working hours of full-time workers, the number of work days, and the implied hourly wages.

During a maternity leave, as an employee is not paid by for any hours by her employer but is instead paid by the Social Security (and may receive a top-up payment from her employer), hours paid are equal to 0. Workers not paid by the hour are an exception to this rule because their hours are imputed based on days paid, which do not vary during maternity leave. As a result, the DADS dataset overestimates hours paid – and underestimates hourly wages – for such workers during years when they give birth to children. In general, these are qualified workers in the upper part of the hourly wage distribution.

Parental earnings and labor supply

I take advantage of a longitudinal version of the DADS dataset that contains detailed information at the person-year level; individuals are identified by an anonymized

personal identifier based on their social security number (NIR), that allows me to track them over time. Starting in 1967, the sample covers individuals born in October in even-numbered years; as of 2002, it also covers individuals born on January 2-5, April 1-4, July 1-4 and October 1-4 regardless of their year of birth. While information on earnings has been available since the creation of the dataset, information on hours paid is only available from 1995, with the exception of central government civil servants, for whom this information is not given before 2009.²⁵

The limitation of this dataset most relevant for my study is the absence of information on self-employment. This may prove problematic if, as suggested by [Connelly \(1992\)](#), mothers tend to turn to self-employment as a more effective way to deal with child-related time constraints. Specifically, if increased affordable childcare provision gives rise to transitions from non-employment and salaried employment towards self-employment, then my estimates will be biased downwards. Conversely, if these increases cause mothers to shift from self-employment to non-employment or salaried employment, then my estimates will be biased upwards. However, this potential bias is likely to be limited: in 2007, less than 5% of mothers with children aged below 3 who held a job were self-employed, and just 4% of those who interrupted their careers were previously self-employed ([Galtier, 2011](#)).

A caveat of the data is that residence is not observed when individuals are not in the labor force. As a result, I have to impute the municipality of residence for person-year observations that correspond to individuals without a job. Specifically, I first impute the municipality of residence to individuals without a job using the municipality of residence when they last held a job, and the municipality of residence when they first held a job for observations that precede the first job held (see [Figure 2.B.1](#)). As a robustness check, I consider the reverse method of imputation (i.e. using the municipality of residence at the time they hold their next job), and find that it does not change my results.

²⁵For these observations, I use a measure of working time expressed in full-time equivalent to impute working hours before 2009. Specifically, the dataset contains information on working time measured in full-time equivalent, that takes values between 0 and 1. I rescale this measure by multiplying it by the median hours of paid work by full-time workers over a full year (1820 hours a year throughout my period of interest).

Individualized childcare

As of 2009, the DADS files also cover salaried employees who are directly paid by households, so this gives me information on childminders and nannies who provide at-home childcare. Specifically, I rely on this feature to compute aggregate earnings and hours paid to childminders and nannies, based on 4-digit occupation,²⁶ at the municipality level from 2009 to 2015, from a comprehensive, cross-sectional version of the dataset.

²⁶The DADS dataset contains an occupation variable based on the *Professions et Catégories Socioprofessionnelles* (PCS) classification. Specifically, I use the most detailed level of this classification, and focus on observations that belong to the "563a – Childminders, baby-sitters and foster families" category.

Figure 2.B.1 – Imputation of the municipality of residence for jobless observations

	Municipality of residence								
Raw data	.	.	A	A	.	.	B	B	.
Baseline imputation	A	A	A	A	A	A	B	B	B
Robustness check	A	A	A	A	B	B	B	B	B

2.B.3 Fertility data

Cross-sectional records

I take advantage of cross-sectional comprehensive birth records to retrieve the mother's municipality of residence at the time of the birth. I use this feature to compute the number of children born to women living in a given municipality in any year between 2005 and 2015. This allows me to recover an approximate measure of the trends in potential demand for childcare at a narrow geographical level.

Longitudinal records

I use on a longitudinal version of these records at the individual level extracted from the *Échantillon Démographique Permanent* (permanent demographic sample, EDP) to obtain information on the timing of births. Thanks to the NIR, this dataset can be merged with the longitudinal version of the DADS. It covers individuals born on October 1-4 whatever their year of birth; information regarding individuals born on January 2-5, April 1-4 and July 1-4 is available from 2004.

A potential issue with the data is the absence of information on children born before 2004 to individuals who were born in January, April and July.²⁷ As a result, the number of children for these individuals is biased downwards. Because my period of interest is 2007-2015, this does not affect the identification of parents of young children in the data, but only the information regarding their past fertility decisions, i.e. whether they have older children or not. For this reason, when controlling for past fertility decisions, I always interact the number of children with a dummy variable that indicates whether the parents are born in October (in which case the data are correct) or not (in which case I underestimate the number of older children that parents have).

²⁷In addition, some birth-related data for the 1990s were incomplete in administrative birth records for individuals born on October 2-3 (for details, see [Wilner, 2016](#)). For these individuals I use the census rather than birth records, as do [Pora and Wilner \(2019\)](#). The quality of these data is comparable to that concerning individuals born on October 1 or 4 for whom administrative birth records are available from 1967.

2.B.4 Sample definition

I exclude all individuals who have ever worked in the childcare industry, so that my results are not driven by the increasing labor demand in this sector. To insure against measurement error in the upper tail of the earnings distribution, and for very low working times, I winsorize earnings at the quantile of order 0.9999, and drop person-year observations that either (i) have fewer than 18 annual paid working days or (ii) have paid daily working hours below 1/20 of legal full-time hours or (iii) have hourly wages under 90% of the minimum wage. For these observations, I consider individuals to be out of the labor force, so that their labor earnings equal 0.

2.C Treatment groups composition

My empirical framework splits municipalities into four groups according to the magnitude of the largest increase in the provision of collective childcare between two consecutive years: a never treated group in which the supply never increases, and then 3 groups based on cutoffs at the 80th and the 90th percentile of the distribution. While my strategy is not based on between groups comparison, in Table 2.C.1, I describe the composition of these groups in terms of municipality characteristics before my time-period of interest begin, i.e. in terms of observables in the 2006 Census.²⁸ This description is nevertheless useful to get a better sense of the composition of the P90-P100 group which is key to my identification strategy.

Specifically, I consider several potentially relevant dimensions:

- municipality size, as defined by the total 2006 population in the census;
- potential and actual birth rates, approximated by:
 - the share of women aged 15 to 49 in the total population;
 - the ratio of children aged less than 1 over the number of women aged 15 to 49;
- migration, as measured by the share of inhabitants who lived, 5 years after the data was collected, in another municipality, either in metropolitan France or abroad;
- couple formation, as measured by the share of single women (men) in the population of women (men) aged 20 to 49;
- marriage formation and dissolution, as measured by:
 - the share of married women (men) in the population of women (men) aged 20 to 49;
 - the share of divorced women (men) in the population of women (men) aged 20 to 49;

²⁸This information was not necessarily collected in 2006: since 2004, the French Census is collected annually. Specifically, all French municipalities are surveyed over a five-year period. As a result, five census surveys were conducted from 2004 to 2008, and were finally combined to produce the Census results, dated 2006, i.e. the medium year. As a result, a slight part of the results relate to a time-period possibly affected by the treatment (year 2008).

- female labor force participation, as measured by the share of women who declared themselves to be housewives in the population of women aged 20 to 49;
- labor market composition, as measured by the share of women (men) who are managers or professionals in the population of women (men) aged 20 to 49.

Because I estimate labor supply effects at the individual-level of parents with potentially affected children, I weight this municipality-level data by the number of children aged 2 or less in 2007. As a result, larger municipalities are given much more weight than smaller municipalities in Table 2.C.

Overall, the P90-P100 group is composed of relatively small municipalities, around 7,600 inhabitants. However, these municipalities do not depart much from other municipalities in terms of potential and actual birth rates, the share of inhabitants who did not live in these municipalities five years before the data was collected or female labor force participation. Marriage rates may be relatively high, and the share of both men and women with managerial or professional positions relatively low.

Table 2.C.1 – Summary statistics at the municipality level: by treatment group

	Never		P25-P80		P80-P90		P90-P100	
# Municipalities	31,205		1,858		763		2,372	
	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation
<i>Collective childcare supply in 2007 (in %)</i>								
Childcare	0.2	5.4	18.2	8.4	20.2	11.4	17.6	23.4
<i>Municipal population (2006)</i>								
Pop.	1,400	1,400	254,200	568,500	21,900	18,500	7,600	7,900
<i>Potential and actual birth rate (in %)</i>								
Pot. mothers	21.9	2.5	25.3	2.6	23.9	2.3	22.9	2.7
Birth rate	5.8	2.5	5.2	1.0	5.3	1.0	5.2	1.3
<i>Immigration rate (in %)</i>								
Mig. (from Fr.)	24.3	7.7	21.1	5.5	23.3	4.8	24.5	6.2
Mig. (abroad)	0.8	1.6	2.4	1.5	1.8	1.1	1.4	1.6
<i>Share of single individuals (in %)</i>								
Single (f)	20.6	6.8	40.5	8.9	34.5	7.5	29.2	8.1
Single (m)	29.8	7.3	43.2	7.7	38.3	6.6	34.9	7.6
<i>Share of married and divorced individuals (in %)</i>								
Married (f)	56.8	7.9	41.4	8.9	46.8	7.3	50.8	7.9
Married (m)	48.6	8.0	37.7	7.8	42.3	6.6	44.9	7.5
Divorced (f)	5.9	2.8	7.8	1.7	7.9	1.8	7.5	2.1
Divorced (m)	4.9	2.4	5.1	1.1	5.2	1.2	5.3	1.5
<i>Female labor force participation (in %)</i>								
Housewives	9.7	5.8	10.7	5.2	9.9	4.6	9.6	4.5
<i>Labor market composition (in %)</i>								
Man. and Prof. (f)	5.5	5.4	10.3	7.6	11.1	8.2	8.4	6.4
Man. and Prof. (m)	9.6	7.9	16.1	10.0	18.3	12.3	14.7	10.3

Note. Data regarding the Tarn département are omitted. *Source.* EAJE-PSU records, CNAF. Birth records and 2006 Census, Insee.

2.D Policy evaluation

2.D.1 Aggregate labor supply effect

Figure 2.5 shows that the evolution of relative childcare supply is extremely flat both before and after the childcare shock. The contribution of municipalities to the national-level increase in affordable collective childcare provision is thus mostly attributable to these shocks, rather than a continuous increase in local supply. Figure 2.D.1, decomposes the yearly growth at national level between (i) the contribution of shocks for each treatment group; (ii) the residual contribution of within-municipality growth when municipalities do not experience a shock; and (iii) a composition shift.²⁹ It clearly shows that the national-level increase is attributable first and foremost to these shocks, especially of those at top of the distribution. As a result, it is largely sufficient to analyze the consequences of these shocks to capture the overall consequences of collective childcare expansions at the national level, and thus to evaluate the impact of the national *plans crèches* over the relevant time period.

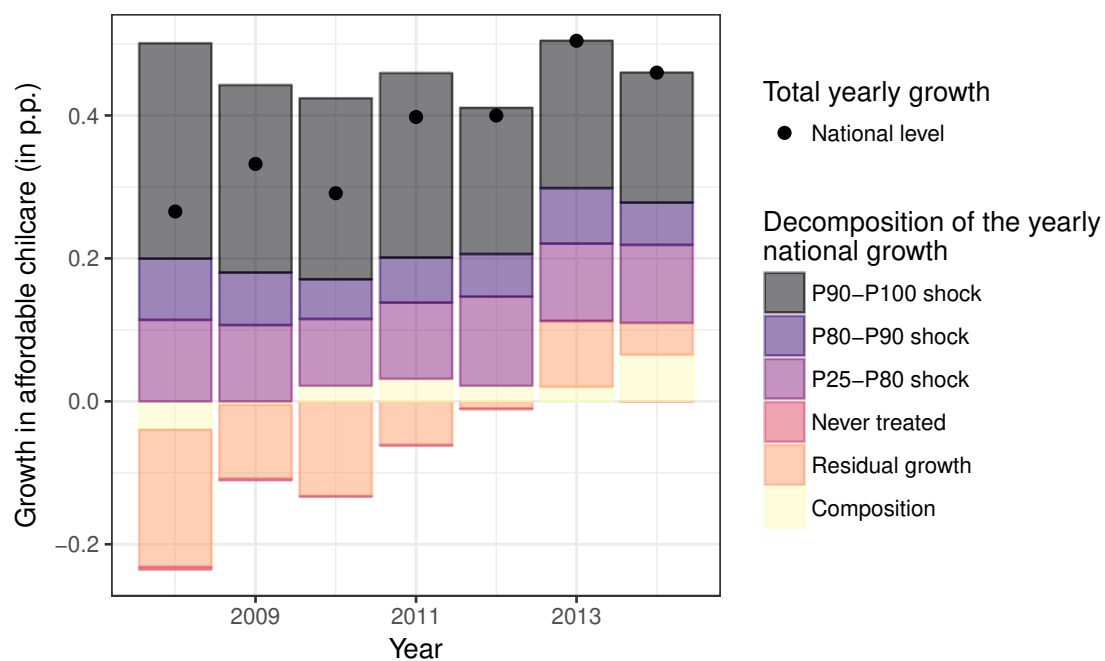
Between 2000 and 2016, 150 000 new childcare places were created under the national plans to expand affordable childcare provision (IGAS/IGF, 2017). Hence, taking my most optimistic estimate from Table 2.4, i.e. the upper bound of my 95% confidence interval, these newly created places enables 8 000 more mothers of young children to hold a salaried job in 2016.

There were about 1.9 million mothers with children under age 3 in 2016, with a salaried employment rate of 67%. As a result, a rough calculation, based on the upper bound of the 95% confidence interval, suggests that the national *plans crèches* contributed, at most, to a 0.4 percentage-point increase in the labor force participation of mothers between 2000 and 2016.³⁰ As a comparison, had the efficiency of the plans reached that of the Norwegian policies investigated by Andresen and Havnes (2019), the counterfactual increase would have amounted to 2.5 percentage points. Empirically, mothers' salaried employment varied very little between 2007 and 2015, fluctuating between 66% and 67%.

²⁹Appendix 2.E details this straightforward accounting decomposition.

³⁰By contrast, my point-estimate estimate implies a 0.1 percentage point drop in the salaried employment rate.

Figure 2.D.1 – Decomposition of national-level annual growth in relative supply of EAJE-PSU affordable collective childcare from 2008 to 2014



Estimates of the contribution of childcare shocks and composition shifts to the annual growth in EAJE-PSU childcare supply at national level (see Appendix 2.E).

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records, CNAF. Birth records, Insee.

2.D.2 Impact on public finances

I map my empirical results into a simple cost-effectiveness evaluation exercise. In the French context where collective childcare is only one among many subsidized childcare solutions, the analysis has to take into account the counterfactual burden for public finances of the childcare solution families would have chosen had they not been offered a place in an EAJE-PSU facility. Because of this, the national plans to expand collective childcare provision may prove beneficial, even with null effects on parental labor earnings and labor supply.

A severe limitation of my data is that they do not allow me to disentangle childminders from at-home childcare provided by nannies, two childcare solutions that have very different consequences for public finances: broadly speaking, the cost for public finances of a place offered by a childminder is 60% that of an EAJE-PSU place, whereas the public cost of tax credits that subsidize at-home childcare represents 120% of an EAJE-PSU place (Figure 2.1).

I design my cost-effectiveness analysis to take into account these crowding-out effects, and the differential burden associated with substitution across childcare solutions. Based on the medium-run effect of my event-study analysis (see Figure 2.8), I assume a total crowding-out of individualized childcare solutions, i.e. that every family who obtained a collective childcare place would have otherwise resorted to either a childminder or at-home childcare. I also assume a null effect as to parental earnings and parental labor supply. As I cannot disentangle one from the other, I consider two cases that can be regarded as bounds on the plausible actual scenario: firstly, if the crowding-out only affects childminders, and secondly, if it only affects at-home childcare.

In the first case, where the substitution effects only affect childminders, the burden for public finances increases. Assuming (i) the annual operating cost of a collective childcare place to be €15,000, of which 82% is covered by public finances, (ii) the collective cost of a childminder place to be 60% of this burden, and (iii) null effects on parental labor supply and total crowding-out, then the long-run burden of creating 150 000 collective childcare places amounts to an additional €738 million of public spending per year. In the second case, i.e. assuming that the substitution effects only affect at-home childcare, and that the public cost of tax credits for this solution to represent 120% those of a collective childcare place, the long-run effect of this plan corresponds to a €369 million reduction in annual public spending.

Table 2.D.1 sums up these results, also considering an alternate counterfactual scenario in which families who receive a childcare place under the national plans are drawn randomly from the population of families that use either a childminder or at-home childcare.³¹ Interpretation of the implied effectiveness-cost ratios is not straightforward as it depends on the directions of two different effects. Firstly, it depends on whether the plans have increased or decreased mothers' salaried employment and salaried earnings: as my confidence intervals always include 0, the worst-case scenario implies a reduction in mothers' labor outcomes. While at first glance, negative maternal labor supply effects seem unlikely, they should not be overlooked, especially since families would use more costly and more flexible childcare solutions had they not been granted a collective childcare place. These counterfactual solutions may therefore offer a better work-life balance than my treatment of interest – e.g. because childminders' working hours are more flexible than those of EAJE-PSU facilities – and the decrease in childcare prices also leaves room for possible income effects.

Secondly, these estimates also depend on whether the national plans produce an increase or a decrease in public spending: substitution from at-home to collective childcare would reduce the burden on public finances. In other words, negative (positive) estimates in the best-case scenario imply both a reduction (increase) in public spending and better labor outcomes for mothers, whereas negative (positive) estimates in the worst-case scenario correspond to an increase (reduction) in public spending paired with poorer labor outcomes for mothers.

As substitution effects lower the impact of the plans on public finances, my estimates appear less precise than they were when omitting these crowding-out effects. The sensitivity of these estimates, and of the implied effectiveness-cost ratios, to assumptions about the composition of the crowding-out effects is extremely salient: these assumptions imply very different conclusions as to the impact of these plans on public finances, not counting the fixed-costs of increasing collective childcare provision. Additional data on individualized childcare would thus be very useful for implementing a full policy analysis of these national plans.

³¹Using a childminder is more than 20 times more frequent than using at-home childcare.

Table 2.D.1 – Empirical policy evaluation: counterfactual scenarios

Counterfactual substitution	No solution	Childminders	At-home childcare	Mixed-case
<i>Long-run operating cost*</i>				
For one place (in €)	+12 300	+4 900	-2 500	+4 600
National plans (in M€)	+1 845	+738	-369	+688
<i>Effectiveness-cost ratio: additional years without career interruption per M€**</i>				
Best case	+4.3	+10.8	-21.7	+11.5
Baseline	-1.4	-3.5	+6.9	-3.8
Worst case	-7.1	-17.8	+35.7	-19.2
<i>Effectiveness-cost ratio: salaried earnings gains per €***</i>				
Best case	+0.18	+0.44	-0.87	+0.47
Baseline	+0.03	+0.08	-0.16	+0.09
Worst case	-0.11	-0.27	+0.53	-0.29

*Excluding the fixed cost paid to create additional childcare places. **The baseline is based on the point estimate, and the best (worst) case scenario is based on the upper (lower) bound of the 95% confidence interval of the estimated employment effect (see Table 2.4). ***The baseline is based on the point estimate, and the best (worst) case scenario is based on the upper (lower) bound of the 95% confidence interval of the estimated salaried earnings effect (see Table 2.4). *Note.* Data regarding the Tarn département are omitted. *Source.* EAJE-PSU records and PAJE records, CNAF. Birth records, comprehensive DADS records and DADS-EDP panel, Insee.

2.E Decomposition of the national annual growth in childcare coverage

Let S_t denote the relative supply of EAJE-PSU affordable collective childcare at the national level on year t :

$$S_t = \frac{N_t^{\text{places}}}{N_t^{\text{birth}} + N_{t-1}^{\text{birth}} + N_{t-2}^{\text{birth}}} \quad (2.6)$$

where N_t^{places} denotes the number of childcare places available at the national level, and N_t^{birth} the number of births that occurred at time t . The national-level supply S_t is a weighted sum of municipality-level supplies, with weights equal to the share of children aged 2 or less who live in each municipality:

$$S_t = \sum_c w_{c,t} S_{c,t} \quad (2.7)$$

As a result, the annual growth in childcare coverage at the national level can be decomposed:

$$\begin{aligned} S_t - S_{t-1} &= \sum_c (w_{c,t} S_{c,t} - w_{c,t-1} S_{c,t-1}) \\ &= \sum_c ((w_{c,t} - w_{c,t-1}) S_{c,t-1} + w_{c,t} (S_{c,t} - S_{c,t-1})) \\ &= \underbrace{\sum_c (w_{c,t} - w_{c,t-1}) S_{c,t-1}}_{\text{Composition}} \\ &\quad + \underbrace{\sum_g \sum_c \mathbb{1}\{c \in g\} \mathbb{1}\{E_c = t\} w_{c,t} (S_{c,t} - S_{c,t-1})}_{\text{Shocks of the treatment group } g} \\ &\quad + \underbrace{\sum_c \mathbb{1}\{E_c \neq t\} w_{c,t} (S_{c,t} - S_{c,t-1})}_{\text{Residual growth}} \end{aligned} \quad (2.8)$$

where g denotes treatment group and E_c denotes the timing of the childcare shock in municipality c . The composition term corresponds to a compositional shift whereby municipalities with higher past childcare supply may expand more quickly or slowly

than their counterparts with lower past coverage. The two other terms correspond to (i) the contribution of shocks for each treatment group; (ii) the contribution of within-municipality growth before or after shocks.

2.F Identification

2.F.1 Sample selection

I develop a very simple framework to provide proof that my sample selection does not impede the identification of the causal effect of childcare on maternal labor supply. Because the municipality of residence is only observed for individuals who are salaried employees, the data only cover individuals who have been, at some point in their lives, salaried employees. Salaried employment being a key measure of labor supply, this raises concerns regarding the validity of my results.

To simplify the problem, I consider a two-period version of the problem: individuals are observed twice, once before childbirth ($t = 0$), and once when they have very young children ($t = 1$). My dependent variable, denoted as Y_{it} , is a dummy variable that equals 1 if individual i is observed in salaried employment at time t , and 0 otherwise. Individuals are assigned a treatment, i.e. a childcare place: this treatment is represented by D_i , a dummy variable. I rely on a potential outcome framework that allows heterogeneous effects of the treatment. In other words, each individual i is associated with a quadruplet $(Y_{it}(d))$ with (t, d) in $\{0, 1\}^2$. The actual outcome writes $Y_{it} = (1 - D_i)Y_{it}(0) + D_iY_{it}(1)$. In the end, individuals are only observed if $Y_{i0} = 1$ or $Y_{i1} = 1$.

I then consider a few simplifying assumptions:

Assumption 1 (Time-monotonicity). $\forall i \in \mathcal{I} \forall d \in \{0, 1\} Y_{i0}(d) \geq Y_{i1}(d)$

Assumption 2 (Exogeneity). $D_i \perp\!\!\!\perp (Y_{it}(d))_{(t,d) \in \{0,1\}^2}$

Assumption 3 (Treatment-monotonicity). $\forall i \in \mathcal{I} Y_{i1}(0) \leq Y_{i1}(1)$

Assumption 4 (No anticipation). $\forall i \in \mathcal{I} Y_{i0}(0) = Y_{i1}(0)$

Assumption 1 states that having children results in women leaving the labor force, or staying in employment, but never in women entering the labor force if they were not employed before; this assumption is based on the large literature devoted to the effect of fertility on labor supply and earnings. Assumption 2 simplifies the research design and considers it to replicate a randomized treatment. Assumption 3 states that being offered a childcare place cannot induce women to leave the labor force, but can only maintain or increase their labor force attachment with respect to the counterfactual situation. Lastly, Assumption 4 states that women do not base

their pre-children labor force participation decision on the future attribution of a childcare place.

These restrictions allow to characterize observed and unobserved individuals in terms of their potential outcomes instead of their realized outcomes:

Observed individuals If $Y_{i0} = 1$ or $Y_{i1} = 1$ then (Y_{i0}, Y_{i1}) belongs to the subset $\{(0, 1), (1, 0), (1, 1)\}$. However, $(0, 1)$ is not allowed by A1. As a result, $Y_{i0} = 1$, so that by A4 $Y_{i0}(0) = Y_{i0}(1) = 1$. Conversely, if $Y_{i0}(0) = 1$ then $Y_{i0} = 1$.

Unobserved individuals If $D_i = 1$, then $Y_{i0} = Y_{i1} = 0$ implies $Y_{i0}(1) = 0$ and $Y_{i1}(1) = 0$, so that by A4 $Y_{i0}(1) = 0$ and by A3 $Y_{i1}(0) = 0$. If $D_i = 0$, then $Y_{i0} = Y_{i1} = 0$ implies $Y_{i0}(0) = 0$ and $Y_{i1}(0) = 0$, so that by A4 $Y_{i0}(1) = 0$, which implies by A1 $Y_{i1}(1) = 0$. Conversely if for all t and all d $Y_{it}(d) = 0$ then $Y_{i0} = Y_{i1} = 0$.

This observation justifies the following result:

Proposition 1 (Average treatment effects). *Under Assumptions 1-4, (i) the difference in average realized outcome between the treatment and the control group among the observed population identifies a local average treatment effect; (ii) the average treatment effect in the overall population equals this estimand multiplied by the share of the observed population.*

Proof. Let us first consider (i):

$$\begin{aligned}
& \mathbb{E}[Y_{i1}|D = 1, Y_{i0} + Y_{i1} > 0] - \mathbb{E}[Y_{i1}|D = 0, Y_{i0} + Y_{i1} > 0] \\
&= \mathbb{E}[Y_{i1}(1)|D = 1, Y_{i0}(0) = 1] - \mathbb{E}[Y_{i1}(0)|D = 0, Y_{i0}(0) = 1] \\
&\stackrel{A2}{=} \mathbb{E}[Y_{i1}(1)|Y_{i0}(0) = 1] - \mathbb{E}[Y_{i1}(0)|Y_{i0}(0) = 1] \\
&= \mathbb{E}[Y_{i1}(1) - Y_{i1}(0)|Y_{i0}(0) = 1]
\end{aligned} \tag{2.9}$$

The average treatment effect in the overall population writes:

$$\begin{aligned}
& \mathbb{E}[Y_i 1(1) - Y_{i1}(0)] \\
= & \mathbb{P}(Y_{i0} + Y_{i1} > 0) \mathbb{E}[Y_i 1(1) - Y_{i1}(0) | Y_{i0} + Y_{i1} > 0] \\
& + \mathbb{P}(Y_{i0} + Y_{i1} = 0) \mathbb{E}[Y_i 1(1) - Y_{i1}(0) | Y_{i0} + Y_{i1} = 0] \\
= & \mathbb{P}(Y_{i0}(0) = 1) \mathbb{E}[Y_i 1(1) - Y_{i1}(0) | Y_{i0}(0) = 1] \\
& + P(\forall t \forall d Y_{it}(d) = 0) \mathbb{E}[Y_i 1(1) - Y_{i1}(0) | \forall t \forall d Y_{it}(d) = 0] \\
= & \mathbb{P}(Y_{i0} = 1) (\mathbb{E}[Y_{i1} | D = 1, Y_{i0} + Y_{i1} > 0] - \mathbb{E}[Y_{i1} | D = 0, Y_{i0} + Y_{i1} > 0]) \quad (2.10)
\end{aligned}$$

□

2.F.2 Childcare shock exogeneity

My event-study framework rests on the assumption that the counterfactual trend in parents' labor outcomes *absent* the local childcare shock is mean-independent of the year when this shock took place. In other words, the decision to open a large number of childcare places should not depend on observables (for decision-makers and not necessarily for the econometrician) that predict different trend in parental labor outcomes before the decision is made.

As I detail in Subsection 2.2.3 and Appendix 2.A.2, the attribution of subsidies directed towards the opening of new childcare places by local CAF offices is mostly based on local measures of childcare coverage, which seems compatible with this assumption. However, municipalities decision to first apply remains unknown, which could seriously question the validity of this assumption.

To assess the plausibility of my identifying assumption in this context, I resort to the 2006 Census at the municipality level. This allows me to test whether, within treatment groups, the timing of childcare expansions correlates with municipality-level characteristics, observed before the decision is made, that could plausibly affect the evolution of parental labor outcomes. Specifically, I consider several potentially relevant dimensions:

- municipality size, as defined by the total 2006 population in the census;
- potential and actual birth rates, approximated by:
 - the share of women aged 15 to 49 in the total population;
 - the ratio of children aged less than 1 over the number of women aged 15 to 49;
- migration, as measured by the share of inhabitants who lived, 5 years after the data was collected, in another municipality, either in metropolitan France or abroad;
- couple formation, as measured by the share of single women (men) in the population of women (men) aged 20 to 49;
- marriage formation and dissolution, as measured by:
 - the share of married women (men) in the population of women (men) aged 20 to 49;

-
- the share of divorced women (men) in the population of women (men) aged 20 to 49;
 - female labor force participation, as measured by the share of women who declared themselves to be housewives in the population of women aged 20 to 49;
 - labor market composition, as measured by the share of women (men) who are managers or professionals in the population of women (men) aged 20 to 49.

This information is extracted from the 2006 Census. However, it was not necessarily collected in 2006: since 2004, the French Census is collected annually. Specifically, all French municipalities are surveyed over a five-year period. As a result, five census surveys were conducted from 2004 to 2008, and were finally combined to produce the Census results, dated 2006, i.e. the medium year. As a result, a slight part of the results relate to a time-period possibly affected by the treatment (year 2008). My results are nevertheless robust to omitting municipalities in which childcare expansions take place in 2008.

I use this information, as well as the level of relative collective childcare supply as measured in 2007, to predict the timing of the municipality-level childcare shock. Specifically, I estimate, separately for each treatment group g , a simple linear model:

$$E_c = \eta_g + \theta'_g X_c + v_c \quad (2.11)$$

where E_c is the date at which municipality c experiences the childcare expansion, η_g a group-specific intercept, X_c a vector of observable characteristics as measured in the 2006 Census, and v_c an idiosyncratic shock of mean 0. To be consistent with the fact that my empirical framework estimates labor supply effects at the individual level of parents with children in the relevant age groups, I weight observations by the number of children aged 2 or less in 2007 (as observed in birth records).

Table 2.F.1 displays my estimates of the vector of coefficients θ_g for the P90-P100 group, which is the most relevant group in my framework.³² It makes it very clear that within the P90-P100 treatment group, municipalities that were treated in the beginning of the 2007-2015 time-period are virtually indistinguishable from their counterparts that were treated later on.

³²Results for the other treatment groups are available upon request.

First, most variables are not significantly correlated with the timing of the childcare expansion. The only variables for which the correlation is significantly different from 0 at usual thresholds are the initial level of relative childcare supply, the city size and the birth rate.

Second, even for these variables, the effect sizes remain tiny: the coefficients would imply for instance that municipalities with initial collective childcare supply 20 percentage points above the mean are treated 0.1 years later in average. Similarly, municipalities with 50,000 inhabitants more than the mean are treated one year earlier: there were only 122 municipalities with population larger than 50,000 in France in 2006. The coefficient on the birth rate may seem large, but the average birth rate in the P90-P100 is 0.05, and its standard deviation is 0.01, which implies very small differences across municipalities.

Third and lastly, even in the full specification, observable characteristics explain very little of the dispersion in the timing of childcare expansions. Indeed, this linear model explains less than 2% of the variance of the timing variable. The poor predictive performance of this models suggests that the exact timing of childcare expansion is too a large extent independent of these characteristics that could plausibly imply different trends in parental labor supply decisions. As a result, it supports the credibility of my key identifying assumption (parallel trends in the counterfactual parental labor supply decisions *absent* the shock).

Table 2.F.1 – OLS estimates of the association between observable characteristics in the 2006 Census and the timing of the local childcare expansion

	(1)	(2)	(3)	(4)	(5)
Childcare	0.56 (0.19)	0.64 (0.19)	0.60 (0.19)	0.60 (0.19)	0.65 (0.20)
Pop. (10,000s)		-0.24 (0.06)	-0.25 (0.07)	-0.26 (0.08)	-0.23 (0.08)
Pot. mothers		-2.68 (1.86)	-3.75 (2.00)	-2.50 (2.17)	-2.33 (2.18)
Birth rate		9.44 (3.53)	9.23 (3.56)	10.24 (3.80)	8.56 (3.96)
Mig. (from Fr.)			0.53 (0.79)	0.68 (0.81)	1.08 (0.86)
Mig. (abroad)			4.59 (3.02)	3.94 (3.16)	3.99 (3.28)
Single (f)				1.01 (2.40)	0.43 (2.47)
Single (m)				-0.13 (2.36)	0.03 (2.37)
Married (f)				2.72 (3.24)	1.84 (3.30)
Married (m)				-1.90 (3.31)	-1.24 (3.35)
Divorced (f)				1.40 (3.07)	0.68 (3.11)
Divorced (m)				5.49 (3.79)	4.94 (3.81)
Housewives					1.38 (1.19)
Managers (f)					-0.19 (1.59)
Managers (m)					-0.30 (1.05)
Observations	2,372	2,372	2,372	2,372	2,372
R ²	0.004	0.02	0.02	0.02	0.02
Adjusted R ²	0.003	0.02	0.02	0.02	0.02

Dependent variable. Timing of the municipality-level childcare shock. *Explanatory variables.* Relative childcare supply as measured in 2007 and observable characteristics at the municipality-level. *Note.* Data regarding the Tarn département are omitted. *Source.* EAJE-PSU records, CNAF. Birth records and 2006 Census, Insee.

2.F.3 Parental earnings

Measurement error

For the individual-level data, the most prominent source of measurement error in my approach stems from the fact that while the treatment assignment is based on the municipality of residence, I do not observe this location when individuals are not in salaried employment. As a result, I choose to impute the location based on past locations, when individuals were salaried employees. This basically makes the strong assumption that individuals who are outside the labor force do not move until they find a new job.

As a robustness check, I make the reverse choice in terms of imputation strategy, i.e. using future municipality of residence to impute location when individuals are outside the labor force. I find my results to be very robust to these changes.

Table 2.F.2 – Instrumental variable estimates of the impact of affordable collective childcare on parents’ labor outcomes, by gender

Age of youngest child	Labor earnings (2015 euros)	Employment (p.p.)	Days	Hours per day	Hourly wages (2015 euros)
<i>Mothers</i>					
0–2	255.7 (860.41)	-2.28 (3.51)	-2.57 (11.23)	-0.066 (0.126)	0.687 (0.472)
<i>Fathers</i>					
0–2	186.29 (1383.8)	1.2 (2.85)	-2.71 (9.4)	0.006 (0.088)	0.052 (0.671)

Dependent variable. Parents’ labor outcomes. *Explanatory variables.* Childcare supply and calendar-time dummies interacted with treatment group, plus municipality fixed effects. Childcare supply is instrumented by time-to-event dummies interacted with treatment group. Standard errors are clustered at the municipality level. *Note.* Data regarding the Tarn département are omitted. *Source.* EAJE-PSU records, CNAF. Birth records and DADS-EDP panel, Insee.

Other policy shocks and compositional shifts

My approach identifies the causal impact of collective childcare on parents' labor outcomes to the extent that the childcare shocks upon which it is based do not correlate with other changes that would affect the outcome. While this assumption cannot be tested directly, it is possible to verify that more restricted versions of this assumption do hold.

To this end, I check that my results are not driven by other local policy changes or shocks to the local labor markets by further interacting the calendar time \times treatment group fixed effects of Model 2.2 with geographical area dummies. As a result, the identification of my parameter of interest stems solely from differences in the timing of the childcare shock across municipalities of the same treatment group and that belong to the same geographical area. I implement this strategy at two distinct levels. First, I consider the *département* level, given that the local offices of the Family branch of the French social security operate at this level. Second, I consider *Zones d'emploi*, a statistical zoning system developed by Insee to delimit local labor markets.³³ This is particularly useful given that my period of interest covers the Great Recession, of which the impact might be heterogeneous across local labor markets. Tables 2.F.3 and 2.F.4 display my results, which are consistent with my main estimates.

To verify that my results are not driven by changes in the composition of potentially treated parents, I also modify Model 2.2 to include individual-level covariates. Specifically, I consider birth cohort (year of birth), education, and past fertility decisions, i.e. total number of children.³⁴ Table 2.F.5 displays my estimates, that are once again consistent with my previous findings.

³³A Zone d'emploi is defined by Insee as a geographical area within which most of the labor force lives and works, and in which firms can find most of the labor force necessary to fill the available jobs.

³⁴In this case, I interact the number of children with a sample dummy (i.e. a dummy variable that indicates whether parents are born on October, in which case their past fertility is perfectly observed, or not, in which case it is left-censored) and calendar time fixed effects to circumvent the left censoring issue mentioned in section 2.3.

Table 2.F.3 – Instrumental variable estimates of the impact of affordable collective childcare on parents’ labor outcomes, by gender

Age of youngest child	Labor earnings (2015 euros)	Employment (p.p.)	Days	Hours per day	Hourly wages (2015 euros)
<i>Mothers</i>					
0–2	191.56 (879.89)	-2.42 (3.53)	-8.39 (11.16)	-0.068 (0.128)	0.731 (0.466)
<i>Fathers</i>					
0–2	418.05 (1333.66)	0.54 (2.7)	8.14 (9.17)	0.007 (0.086)	0.249 (0.642)

Dependent variable. Parents’ labor outcomes. *Explanatory variables.* Childcare supply plus calendar-time dummies interacted with treatment group and departement dummies, plus municipality fixed effects. Childcare supply is instrumented by time-to-event dummies interacted with treatment group. Standard errors are clustered at the municipality level. *Note.* Data regarding the Tarn département are omitted. *Source.* EAJE-PSU records, CNAF. Birth records and DADS-EDP panel, Insee.

Table 2.F.4 – Instrumental variable estimates of the impact of affordable collective childcare on parents' labor outcomes, by gender

Age of youngest child	Labor earnings (2015 euros)	Employment (p.p.)	Days	Hours per day	Hourly wages (2015 euros)
<i>Mothers</i>					
0–2	-256.91 (893.63)	-4.23 (3.44)	-3.12 (11.39)	-0.062 (0.128)	0.889 (0.512)
<i>Fathers</i>					
0–2	879.48 (1498.4)	3.1 (2.82)	6.69 (9.21)	0.052 (0.086)	0.208 (0.733)

Dependent variable. Parents' labor outcomes. *Explanatory variables.* Childcare supply plus calendar-time dummies interacted with treatment group and Zone d'emploi dummies, plus municipality fixed effects. Childcare supply is instrumented by time-to-event dummies interacted with treatment group. Standard errors are clustered at the municipality level. *Note.* Data regarding the Tarn département are omitted. *Source.* EAJE-PSU records, CNAF. Birth records and DADS-EDP panel, Insee.

Table 2.F.5 – Instrumental variable estimates of the impact of affordable collective childcare on parents’ labor outcomes, by gender

Age of youngest child	Labor earnings (2015 euros)	Employment (p.p.)	Days	Hours per day	Hourly wages (2015 euros)
<i>Mothers</i>					
0–2	778.53 (806.14)	-0.49 (3.4)	2.4 (10.93)	0.006 (0.121)	0.706 (0.425)
<i>Fathers</i>					
0–2	567.93 (1333.54)	2.98 (2.8)	-2.92 (9.25)	0.002 (0.087)	-0.05 (0.628)

Dependent variable. Parents’ labor outcomes. *Explanatory variables.* Childcare supply plus calendar-time dummies interacted with treatment group, plus municipality fixed effects and parents’ education interacted with birth cohort (year of birth) and total number of children (interacted with a sample dummy and calendar time dummies). Childcare supply is instrumented by time-to-event dummies interacted with treatment group. Standard errors are clustered at the municipality level. *Note.* Data regarding the Tarn département are omitted. *Source.* EAJE-PSU records, CNAF. Birth records and DADS-EDP panel, Insee.

Placebo groups

However, these checks are not sufficient to assess the credibility of my design if childcare expansions coincide with shocks that take place at the municipality level, as opposed to the département or Zones d'emploi level, and if these shocks do not correlate with compositional shifts in pool of potentially treated parents. As an additional attempt to test the credibility of my identification strategy, I replicate it on placebo groups, i.e. subsets of individuals who should not be directly affected by childcare expansions.

To this end, I consider two groups: parents taken one to five years before the birth of their first child, and parents whose youngest child is aged 3 to 10. Because EAJE-PSU facilities target children aged 0 to 2, there should not be any direct effect of childcare expansions on the labor outcomes of these groups. Results appear in Table 2.F.6. Consistent with this rationale, I cannot detect any significant effect for these groups, which strengthens the credibility of my findings. If anything, the earnings of fathers may even increase in municipalities that experience the largest shocks. This would be the case if these massive expansions were more likely to occur early in municipalities where the labor market is growing steadily. However, this would suggest that my estimates for the labor market outcomes of mothers are biased upwards, which gives even more strength to my claim that the positive maternal labor supply effects of childcare expansions are negligible at best.

Table 2.F.6 – Instrumental variable estimates of the impact of affordable collective childcare on parents' labor outcomes, by gender

Age of youngest child	Labor earnings (2015 euros)	Employment (p.p.)	Days	Hours per day	Hourly wages (2015 euros)
<i>Mothers</i>					
<i>a. Full sample</i>					
-5--1	-710.47 (1065.09)	-3.23 (3.99)	-1.02 (15.72)	-0.008 (0.146)	-0.412 (0.44)
3-10	-27.31 (793.09)	0.04 (2.56)	8.62 (9.1)	0.079 (0.091)	-0.626 (0.393)
<i>b. P90-P100 treatment group</i>					
-5--1	-1008.4 (1116.39)	-3.95 (4.26)	4.07 (16.97)	0.029 (0.156)	-0.503 (0.459)
3-10	552.18 (819.24)	1.65 (2.65)	12.09 (9.53)	0.131 (0.096)	-0.522 (0.404)
<i>Fathers</i>					
<i>a. Full sample</i>					
-5--1	1389.02 (1477.61)	5.76 (3.91)	-1.01 (14.13)	-0.122 (0.123)	-0.349 (0.758)
3-10	2170.21 (1288.64)	2.9 (2.28)	-3.66 (7.9)	0.064 (0.073)	0.757 (0.657)
<i>b. P90-P100 treatment group</i>					
-5--1	1512.79 (1543.95)	4.02 (4.01)	-0.4 (15)	-0.04 (0.131)	-0.104 (0.798)
3-10	2897.72 (1353.67)	3.95 (2.35)	0.18 (8.2)	0.052 (0.076)	0.872 (0.692)

Dependent variable. Parents' labor outcomes. *Explanatory variables.* Childcare supply and calendar-time dummies interacted with treatment group, plus municipality fixed effects. Childcare supply is instrumented by time-to-event dummies interacted with treatment group. Standard errors are clustered at the municipality level. *Note.* Data regarding the Tarn département are omitted. *Source.* EAJE-PSU records, CNAF. Birth records and DADS-EDP panel, Insee.

Sampling issues

It could be argued that because (i) I focus on very modest expansion plans (see Section 2.2); (ii) I rely solely on information on the local aggregate provision of childcare, as opposed to individual offers made to parents; and (iii) I do not rely on comprehensive data but only on a sample of parents, my null effects are driven by the fact that my sample does not actually include any parents who received a recently created childcare place. The probability of this event is not 0, so this is always a possibility. However, with mild assumptions I am able to quantify the probability of these huge deviations.

First, the data indicates that 70 000 childcare places were created between 2007 and 2015 (see Section 2.2). Assuming that (i) this increase was linear overtime (see Figure 2.2) and that (ii) childcare places are reallocated every three years,³⁵ the increase resulted in 105 000 additional childcare allocation decisions. Second, my sample of parents is based on their birthday, i.e. whether or not they were born in the 16 relevant days of the EDP sample (see Section 2.3), so the sampling rate is 4.4%.

Consider $\hat{n}^{\text{allocations}}$ the number of additional childcare allocation decisions stemming from the national increase in coverage that benefit the parents in my sample. For each decision i , let B_i denote a dummy variable that is equal to 1 if this spot is allocated to a parent of my sample. Then:

$$\hat{n}^{\text{allocations}} = \sum_{i=1}^M B_i \quad (2.12)$$

Given that the allocation of childcare places does not depend on whether *parents* are born on the 16 EDP days or not,³⁶ B_i variables can be assumed to be independent Bernoulli variables of parameter p equal to the sampling rate. As a result, the variance of $\hat{n}^{\text{allocations}}$ is written $Mp(1-p)$. It is then easy to apply Chebychev's inequality:

$$\mathbb{P}(|\hat{n}^{\text{allocations}} - Mp| \geq \rho Mp) \leq \frac{1-p}{\rho^2 Mp} \quad (2.13)$$

³⁵This is the most conservative estimate. If childcare places are reallocated more frequently, then the probability of large deviations decreases.

³⁶It may depend on their *children's* date of birth given that childcare places are frequently offered from September to July, when made available by children aged 3 leaving the EAJE facilities to attend preschool.

With $M = 105,000$, $p = 0.044$ and $\rho = 0.1$, an upper bound for the probability that the number of additional childcare allocation decisions that benefit parents in my sample deviates by more than 10% from its expected value is 2%.


Event-study framework

My graphical analysis is based on an event-study framework that relies on a no-pretrend assumption to disentangle calendar-time from time-to-treatment effects. However, this approach is based on the premise that while treatment effects may be dynamic, i.e. they may vary depending on whether units are observed one year, two years, etc. after childcare expansions, treatment effects are actually homogeneous across units that belong to different cohorts, as defined by the timing of the childcare shocks.

In a recent investigation of this setting, [Sun and Abraham \(2020\)](#) show that when this homogeneity assumption fails, the canonical event-study estimates are a weighted sum of cohort-specific treatment effects with many potentially negative weights. As a result, this framework does not enable to properly test the hypothesis that treatment effects are equal to 0 before treatment. Instead, the authors propose to estimate a fully-interacted model, and then to manually average coefficients over cohorts using weights proportional to sample size.

Figure [2.F.1](#) replicates my event-study analysis based on their approach. My results are robust to this concern.

Figure 2.F.1 – Event-study estimates based on [Sun and Abraham \(2020\)](#) of the impact of the childcare shock on mothers' labor earnings, by treatment group



Event-study estimates of the effect of childcare shocks on mothers' labor earnings (Model 2.2).

Source. EAJE-PSU records, CNAF. Birth records and DADS-EDP panel, Insee.

Fuzzy difference-in-difference

The fuzzy difference-in-difference setting proposed by [Duflo \(2001\)](#) which serves as the basis of my approach has been recently subjected to investigation by [de Chaisemartin and DHaultfœuille \(2018\)](#), who raise concerns about its ability to identify causal parameters of interest in realistic settings. They show that under usual assumptions, the canonical Wald-DID estimator only identifies a Local Average Treatment Effect (LATE)³⁷ if either (i) treatment effects are homogeneous; or (ii) the treatment rate (here, the childcare coverage) is constant in the control group.

They propose several corrections that make it possible to identify this LATE when neither of these assumptions are plausible. However, these corrections require either (i) that actual treatment (i.e. the use of a childcare place) be observed at the individual level; or (ii) that the outcome be continuously distributed. In my particular setting, I cannot observe whether individuals do indeed take up collective childcare places, so I cannot use individuals who land a childcare place before childcare expansions to implement their correction. Additionally, my outcomes of interest are not continuously distributed: given that labor supply decisions at the extensive margin are at play, the distribution of labor earnings displays a large mass at 0. It would appear, therefore, that the alternative estimators developed by [de Chaisemartin and DHaultfœuille \(2018\)](#) are not applicable here.

However, this problem is probably not a major threat to my identification strategy. First, the treatment rate varies very little except for the shock, as made evident by [Figure 2.5](#), so that deviations from the assumption that this treatment rate is actually constant are small or non-existent. Second, my approach can be replicated in a setting where, in the control group of municipalities, the childcare coverage rate is constant by construction. To this end, I restrict the analysis to municipalities with no EAJE-PSU facility in 2007, and where a facility opened at some point between 2008 and 2014, and define my childcare shock as the opening of this first facility. In this setting, the treatment rate in the control group, i.e. in municipalities where a collective childcare facility will open at some point, but has not done so yet, is by construction equal to 0.

[Table 2.F.7](#) displays the corresponding Wald-DID estimates. While my standard errors are larger, because I only rely on a very restricted subset of municipalities,

³⁷Specifically: the average treatment effect for those individuals who are offered a childcare place due to the childcare expansion, but would not have been so had they been observed before the expansion.

the results are in line with those obtained using all childcare shocks: I cannot detect any significant change in the labor earnings of mothers with young children after the creation of a childcare facility.

Table 2.F.7 – Instrumental variable estimates of the impact of affordable collective childcare on parents’ labor outcomes based on the opening of the first EAJE-PSU facility, by gender

Age of youngest child	Labor earnings (2015 euros)	Employment (p.p.)	Days	Hours per day	Hourly wages (2015 euros)
<i>Mothers</i>					
0-2	-496.21 (1037.64)	0.49 (4.7)	-20.79 (14.4)	-0.036 (0.151)	0.03 (0.537)
<i>Fathers</i>					
0-2	-1390.23 (1461.44)	1.09 (3.55)	-21.23 (12.96)	0.158 (0.112)	-0.76 (0.806)

Dependent variable. Parents’ labor outcomes. *Explanatory variables.* Childcare supply and calendar-time dummies, plus municipality fixed effects. Childcare supply is instrumented by time-to-event dummies. Standard errors are clustered at the municipality level. *Note.* Data regarding the Tarn département are omitted. *Source.* EAJE-PSU records, CNAF. Birth records and DADS-EDP panel, Insee.

2.F.4 Substitution effects

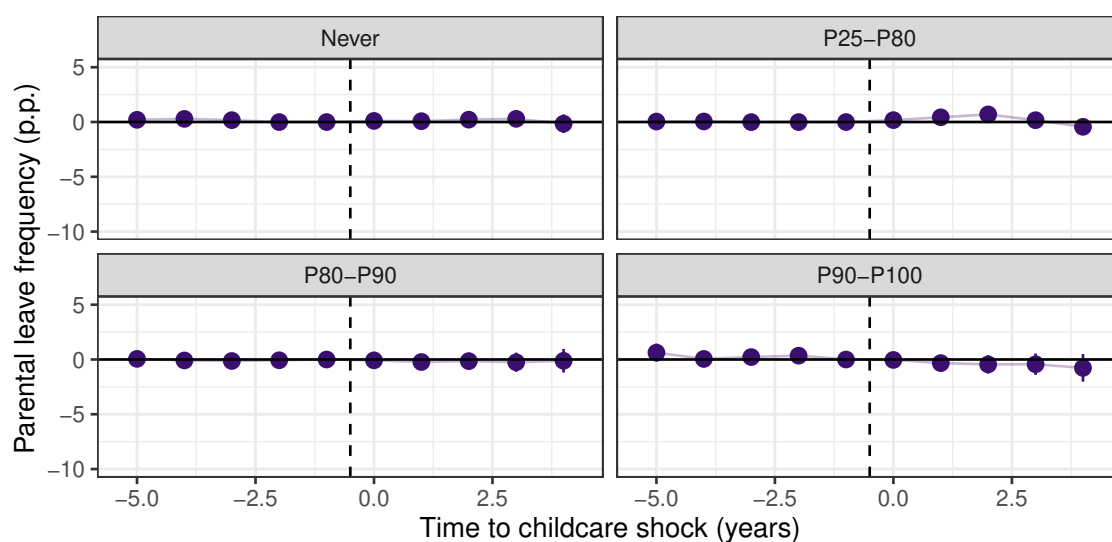
Placebo groups

As what was the case when investigating labor outcomes effects, my division of municipalities in four treatment group allows me to consider effects in municipalities where changes in the supply of collective childcare were actually either non-existent or negligible, which can be regarded as placebo groups. Even though standard errors may be large, I cannot detect any change in demand for individualized childcare in these municipalities, which is reassuring as to the validity of the assumptions upon which my identification strategy is based. This also holds for my estimates regarding the impact of collective childcare on paid parental leave take-up.

Division bias

I then investigate whether these results are affected by some kind of division bias. This could be the case because I use a measure of the number of children aged 2 or less as the denominator in both my measure of collective and individualized childcare and my measure of the frequency of paid parental leave. As a result, measurement error in this number may generate spurious correlations; specifically, the correlation between the supply of different types of childcare solutions might be biased towards 1. To investigate this possibility, I replicate my analysis while adding my measure of the number of children as a covariate in the regression. Figures [2.F.2](#) and [2.F.3](#) display my estimates. I find that this does not change my results.

Figure 2.F.2 – Event-study estimates of the impact of the childcare shock on paid parental leave take-up, by treatment group, controlled for changes in the number of children




Event-study estimates of the effect of childcare shocks on the share of families with children aged 2 or less that receive parental leave allowances in December.

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records and PAJE records, CNAF. Birth records, Insee.

Figure 2.F.3 – Event-study estimates of the impact of the childcare shock on the supply of individualized childcare, by treatment group, controlled for changes in the number of children



Event-study estimates of the effect of childcare shocks on the relative supply of individualized childcare by childminders and nannies.


Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records, CNAF. Birth records and comprehensive DADS records, Insee.

Other policy changes

As I did for labor supply effects, I verify that my results are not driven by other non-relevant policy changes by further conditioning my analysis on various geographical units. Specifically, I consider the *département*, the level at which local offices of the CNAF operate, the *Zone d'emploi*, to consider local labor market effects, and lastly the *Bassin de vie*, a geographical unit defined by Insee that captures the provision of local services at a narrow geographical level. For all these levels, I replicate my analysis while adding treatment group \times calendar time \times geographical unit fixed effects. Figures 2.F.4 to 2.F.9 display the corresponding estimates. I find qualitatively similar effects, which supports the validity of my identification strategy.

Figure 2.F.4 – Event-study estimates of the impact of the childcare shock on paid parental leave take-up, by treatment group, with département-level calendar time fixed effects




Event-study estimates of the effect of childcare shocks on the share of families with children aged 2 or less who received parental leave allowances in December.

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records and PAJE records, CNAF. Birth records, Insee.

Figure 2.F.5 – Event-study estimates of the impact of the childcare shock on paid parental leave take-up, by treatment group, with Zone d’emploi-level calendar time fixed effects




Event-study estimates of the effect of childcare shocks on the share of families with children aged 2 or less who received parental leave allowances in December.

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records and PAJE records, CNAF. Birth records, Insee.

Figure 2.F.6 – Event-study estimates of the impact of the childcare shock on paid parental leave take-up, by treatment group, with Bassin de vie-level calendar time fixed effects




Event-study estimates of the effect of childcare shocks on the share of families with children aged 2 or less who received parental leave allowances in December.

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records and PAJE records, CNAF. Birth records, Insee.

Figure 2.F.7 – Event-study estimates of the impact of the childcare shock on the supply of individualized childcare, by treatment group, with département-level calendar time fixed effects




Event-study estimates of the effect of childcare shocks on the relative supply of individualized childcare by childminders and nannies.

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records, CNAF. Birth records and comprehensive DADS records, Insee.

Figure 2.F.8 – Event-study estimates of the impact of the childcare shock on the supply of individualized childcare, by treatment group, with Zone d’emploi-level calendar time fixed effects




Event-study estimates of the effect of childcare shocks on the relative supply of individualized childcare by childminders and nannies.

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records, CNAF. Birth records and comprehensive DADS records, Insee.

Figure 2.F.9 – Event-study estimates of the impact of the childcare shock on the supply of individualized childcare, by treatment group, with Bassin de vie-level calendar time fixed effects



Event-study estimates of the effect of childcare shocks on the relative supply of individualized childcare by childminders and nannies.

Note. Data regarding the Tarn département are omitted.

Source. EAJE-PSU records, CNAF. Birth records and comprehensive DADS records, Insee.

Chapter 3

Do children explain nurses shortages?

3.1 Introduction

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

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

This paper approaches these issues by putting motherhood at the center of the analysis. Specifically, I quantify how much the causal effect of motherhood on women's labor supply explains of the labor supply lifecycle profile of French hospital nurses. Indeed, while several papers note that the presence of young children correlates with decreases in the labor supply of nurses – either in terms of hours

worked (Askildsen et al., 2003; Hanel et al., 2014) or in terms of participation decision (Phillips, 1995; Nooney et al., 2010) –, very few of them further investigate this fact.¹ By contrast, I build on recent insights from the gender gap literature, according to which most of the differences in labor market outcomes between men and women result from multiple decisions, and especially negative labor supply decisions along multiple margins, that mothers make in response to the arrival of children (Kleven et al., 2019). Nurses being an extremely feminized profession – in France, over 85% of nurses are women (Bessière, 2005) – and recent evidence that the supply of nursing can react strongly to parental leave policies (Friedrich and Hackmann, 2021) suggest that poor family-work conciliation and gender norms regarding child rearing may actually explain a substantial share of the insufficiency of nurses labor supply.

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

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

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

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

The remainder of this paper is organized as follows. Next section outlines the institutional context. Section 3.3 describes the administrative data upon which the analysis is based. Section 3.4 details my empirical framework. Section 3.5 presents the results and Section 3.6 concludes.

3.2 Institutional context

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

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

There are two main settings in which nurses may work in France. Firstly, nurses may work as salaried employees, whose employer may be a hospital, a long-term care

facility, a health center, a school etc. Secondly, they may work as freelance nurses, in which case they provide healthcare directly to patients. However, freelance nursing is not open to fresh nursing graduates. Indeed, being granted the authorization to work as a freelance nurse requires not only a nursing diploma, but also at least two years of experience as a salaried nurse at a healthcare facility. In 2006, the vast majority of registered nurses (63%) were salaried employees at a hospital, either public (49%) or private (14%); 4% of them were employed in a long-term care facility, and 21% of them were salaried employees in other settings. Lastly, 12% of them worked as freelance nurses (Barlet and Cavillon, 2010). In this paper, I focus on the lifecycle of nurses who hold a job at a hospital at one point in their lives. This actually covers the vast majority of the nursing profession, because over three quarters of nurses work at a hospital when they begin their careers (see Appendix 3.B).

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

3.3 Data

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

3.3.1 Labor market data

My labor market data are drawn from the *Déclarations Annuelles de Données Sociales* (DADS). By law,² French employers have to fill in a DADS form for every

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

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

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

These data have two main caveats with respect to my analysis. The first one is that the most detailed occupation variable is not available before 2009. Indeed, before this employers only had to answer a 2-digits occupation question, as opposed to the 4-digits occupation which is the most detailed level in the occupation classification used by Insee. This prohibits the naive approach to the labor supply of nurses, which would basically select individuals into the sample based on whether or not they have, at one point in their lives, held a job as a nurse, as made salient by the 4-digits occupation variable. Instead, I select individuals based on the combination of the 2-digits occupation variable, and the 5-digits industry variable. Subsection 3.3.3 details this choice and characterizes the selected individuals in terms of their detailed occupation, when observed.

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

The second issue is that hours worked are not observed before year 1995. Before this, the data only provide information on days worked, and working-time status, either full-time or part-time. I choose to impute hours worked, measured in full-time units, before 1995, based on those two variables. Specifically, for full-time workers, I consider time worked to be proportional to days worked, and equal to 1 for those who work for a full year – in this there is absolutely no difference with the way time worked is measured after 1995. For part-time workers, I consider them to be on a 50% schedule, which was the most frequent case after 1995 (see Appendix 3.A); as a result, time worked is proportional to days worked, so that for those who work for an entire year time worked is equal to 0.5. My results are nevertheless robust to this particular choice (see Subsection 3.5.3).

3.3.2 Fertility data

My analysis also relies on birth records. Births are registered by an individual who was present at the time of birth, usually the father, but in some cases a doctor or a midwife. I take advantage of a longitudinal version of these records at the individual level extracted from the *Échantillon Démographique Permanent* (permanent demographic sample, EDP) to obtain information on the timing of births. Because it displays an anonymized personal identifier based on the same social security number as the DADS data, this dataset can be merged with the longitudinal version of the DADS.

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

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

3.3.3 Sample construction

My analysis relies on a sample of qualified healthcare workers working in a hospital setting, that I follow over the course of their lives. As explained in Subsection 3.3.1, data regarding the detailed occupation are not available before 2009, which restricts the possibility to base my sample selection on this variable. Instead, I rely on the combination of the 2-digits occupation and 5-digits industry. Specifically, I define hospital nurse jobs as those with (i) 2-digits occupation variable equal to "Qualified healthcare and social workers"⁴ and (ii) 5-digits industry variable equal to "Hospital activities". This definition only approximates the usual definition of a hospital nurse job. However, it matches the usual definition quite closely: Table 3.1 displays the distribution of detailed occupations among jobs that match this criterion, which is observed in the DADS data as of 2009. Over three quarters of these jobs are indeed nurse jobs; the remainder are mostly health related technical jobs. Note that my approach includes nurses managers, but excludes auxiliary nurses who belong to another 2-digit occupation group. Social workers, although nominally included in the criterion, represent a very small share of this population.

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

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

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

⁵Additionally, the vast majority of individuals who ever hold a nurse position will fall within the universe that I cover, because most of them work at a hospital at some point of their lives. To show this, in Appendix 3.B I quantify the share of workers who ever hold a job at a hospital among those who I can properly identify as beginning their careers as nurses after 2010 thanks to the detailed occupation variable. This share is about 75%.

spent outside employment, but they do not take into account the labor supply in the freelance sector.

Crucial to my analysis are (i) the year during which individuals are first observed to hold a nurse job at a hospital and (ii) the length of the time-period during which they are observed afterwards. That my data only cover the 1988-2017 time-period generates two issues with respect to this. Firstly, a substantial share of individuals who are observed to hold such a job in 1988 are likely to have done so for an unobserved number of years, which creates a left-censoring issue. Secondly, recent cohorts, as defined based on the timing of the first hospital nurse job, are only observed for a restricted time-period afterwards. My identification strategy is entirely based on within-cohort comparisons of nurses who become mothers sooner or later. Hence, for these recent cohorts, some control groups of mothers whose first child is born later are unobserved, because the birth of this child cannot be recovered from the data. This creates a right-censoring issue. I discuss these two concerns in Subsection 3.5.3 and provide evidence that my results are robust with respect to these issues.

3.3.4 Summary statistics

Table 3.2 displays a few summary statistics regarding the sample. Individuals of interest usually get their first job as a hospital nurse around age 30, usually after a few years in different jobs, either as a nurse in a different work setting, or in a different occupation. By age 45, about 80% of them have children. This rate is comparable to that of the overall French population (81,7% for women, see [Reynaud, 2020](#)). As in 2015 in France, mothers' average age when they gave birth to their first child was 28.5, this suggests that for a large share of hospital nurses, the beginning of their career coincides with child-rearing years. The remainder of the paper aims at quantifying this fact and its implications.

To this end, Figure 3.1 plots the share of mothers over time relative to the first hospital nurse job, both in the aggregate and separately across cohorts defined by the timing of the first hospital nurse job. The share of mothers increases slightly before nurses land their first job at a hospital, but remains quite small: one year before they get their first job as a qualified healthcare worker at a hospital, less than one in five nurses are mothers. By contrast, after they get their first nurse job at a hospital, this share rises quickly: ten years after the said first nurse job at a

hospital, the share of mothers is 66%. Note that this share is not conditional on gender: among women the proportion is about 80%. For the sake of this particular paper, considering the motherhood rate unconditional on gender has merits because what matters for the aggregate supply of nursing labor is the unconditional rate, as opposed to the fertility rate of women. The difference between the rate at the onset of a career and the rate several years afterwards implies a quick expansion of the number of mothers soon after nurses land their first job at a hospital, as made obvious by the figure. Splitting the data across cohorts confirms that this pattern is not driven by changes in the composition of cohorts that are observed at each point of time, given the restricted time-period of observation.

Figure 3.2 displays the labor supply, that is (a) the average hours worked, measured in full-time units and (b) the salaried employment rate, over time relative to the first hospital nurse job, both in the aggregate and across cohorts defined by the timing of the first hospital nurse job. Soon after the first two years, which correspond to a gradual entry into the job, the labor supply in the salaried sector starts to decay. This decline amounts to 0.16 full-time units over the first ten years of a career, which is substantial considering that the baseline at the beginning of a career is about 0.85 full-time unit. This decline is to a large extent driven by decisions at the extensive margin, i.e. participation decisions, as made obvious by the large decrease in the salaried employment rate. Once again, this pattern holds across cohorts and is not driven by changes in composition over time.

The data allow me to further details these results. Indeed, the DADS data are available at the job spell level, that is at the individual \times employer \times year. As a result, hours worked by an individual at a given point in time can be decomposed as the sum of hours worked across all employers. Specifically, in this paper, I consider three decompositions. Firstly, I contrast hours worked as a nurse with hours worked in other jobs: nurse jobs are those that belong to the "Qualified healthcare and social workers" 2-digit occupational group that I use in the sample definition. Secondly, I compare hours worked at a hospital with hours worked in other work settings: hours worked at hospital is the sum of hours worked across employers that belong to the "Hospital activities" 5-digit industry according to the data. Lastly, I confront hours worked in the public sector with hours worked in the private sector, based on the sector to which employers belong. Because inclusion in my sample is based on whether individuals hold a specific type of job once in their lives, these decompositions are useful to document whether hospital nurses transition to other

types of jobs over the course of their lifecycle.

Figure 3.3 relies on these decompositions to compare the lifecycle profiles of (a) hours worked at a hospital vs. in another setting; (b) hours worked as a nurse vs. in another occupation; and (c) hours worked in the public sector vs. in the private sector. It makes it clear that the decline in labor supply in the salaried sector is first and foremost driven by a decline in hours worked at a hospital, as a nurse, and to a lesser extent in the public sector. Interestingly, although hours worked in non-hospital or non-nurse jobs do slightly increase, the magnitude of this rise remains limited. This implies that the decay in hours worked as a hospital nurse over the lifecycle is not driven by the reallocation of hospital nurses towards other salaried jobs.

3.4 Empirical analysis

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

3.4.1 Model and identification

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

I define $Y_{i,t}(c)$ to be the potential labor supply of individual i at time t had she

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

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

$$Y_{i,t} = Y_{i,t}(\infty) + \sum_c (Y_{i,t}(c) - Y_{i,t}(\infty)) \mathbf{1}\{C_i = c\} \quad (3.1)$$

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

$$CATT_{g,c,t} = \mathbb{E}[Y_{i,t}(c) - Y_{i,t}(\infty) \mid G_i = g, C_i = c] \quad (3.2)$$

This quantity corresponds to the effect of being $t - c$ years away from the birth of one's first child, for those who gave birth to their first child at time c , and belong to group g . These average treatment effects are not conditional on possible subsequent childbirths. As a result, it incorporates both the causal effect of motherhood at the extensive margin, i.e. choosing to be a mother or not, and that of to the intensive margin, i.e. choosing to give birth to one additional child for those who are already with child. In other words, the causal effect of motherhood mixes that of the first child, and of all subsequent children, with weights that depend on the difference between the time-period t and the timing of the first child's birth c : short-run effects ($t = c$) relate almost exclusively to the extensive margin of fertility, whereas longer run effects ($t > c$) will integrate a larger share of the consequences of the intensive margin. This is especially true in a context in which most parents choose to have more than one child, as implied by Table 3.2. I discuss these concern, and provide a decomposition of the effect of children between these two margins in Appendix 3.C.2.

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

Assumption 5 (Parallel trends in baseline outcome). *For all g , for all (t, t') , for all (c, c') , if $c, c' > 1$ and $c, c' < \infty$ then:*

$$\mathbb{E}[Y_{i,t}(\infty) - Y_{i,t'}(\infty) \mid G_i = g, C_i = c] = \mathbb{E}[Y_{i,t}(\infty) - Y_{i,t'}(\infty) \mid G_i = g, C_i = c'] \quad (3.3)$$

Assumption 6 (Limited anticipation). *For all t , for all g , for all c , if $t < c - 1$*

then:

$$\mathbb{E}[Y_{i,t}(c) - Y_{i,t}(\infty) | G_i = g, C_i = c] = 0 \quad (3.4)$$

Assumption 5 states that absent children, the average labor supply of mothers who were born at the same time, took their first job at a hospital during the same year and had their first child only afterwards, would evolve in parallel over time. Assumption 6 states that the average effect of children on their mothers' labor supply is 0 up until two years before they are born. The reason for this choice, as opposed to a full no-anticipation assumption is that (i) becoming a mother during year t generally results from fertility decisions that were made during year $t - 1$, and (ii) maternity leave may start during the last year before childbirth in case childbirth happens at the beginning of the civil year, which will mechanically affect the mother's labor supply.

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

Proposition 2 (Difference-in-difference estimand). *For all (g, c, t) , if $1 < c < \infty$ then:*

$$\begin{aligned} CATT_{g,c,t} &= \mathbb{E}[Y_{i,t} | G_i = g, C_i = c] \\ &\quad - \mathbb{E}[Y_{i,c-2} | G_i = g, C_i = c] \\ &\quad - \mathbb{E}[Y_{i,t} | G_i = g, \max(1, c - 2, t + 1) < C_i < \infty] \\ &\quad + \mathbb{E}[Y_{i,c-2} | G_i = g, \max(1, c - 2, t + 1) < C_i < \infty] \end{aligned} \quad (3.5)$$

Proposition 2 implies that as long as, within a group, there are mothers to be that can still be observed at least two years before their first child is born, it is actually possible to impute the counterfactual labor supply lifecycle profile of mothers whose first child is already born, so as to identify cohort-specific ATTs. Specifically, let $\{\underline{T}(g), \underline{T}(g) + 1, \dots, \overline{T}(g) - 1, \overline{T}(g)\}$ denote the set of time-periods that can be observed for individuals who belong to group g . Then for all c , $CATT(g, c, t)$ is identified from the data provided that:

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

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

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

This last condition implies that very long run effects generally cannot be identified under these assumptions, because no counterfactual is available after the last mother is about to have her first child. Specifically, given the profile of Figure 3.1, cohort-specific effects are very unlikely to be identified for $t \geq 10$, and even more so for $t \geq 15$.

3.4.2 Aggregation and estimation

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

$$\tau(s) = \mathbb{E}[Y_{i,C_i+s}(C_i) - Y_{i,C_i+s}(\infty) \mid (G_i, C_i, C_i + s) \in \Omega] \quad (3.6)$$

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

$$\tau(s) = \sum_{(g,c,c+s) \in \Omega} \mathbb{P}(G_i = g, C_i = c \mid C_i(G_i, C_i, C_i + s) \in \Omega) CATT(g, c, c + s) \quad (3.7)$$

By Proposition 2, it is therefore possible to express $\tau(s)$ as a function of quantities that are all identified from the data.

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

$$\bar{\tau}(S) = \frac{1}{S} \sum_{s=0}^S \tau(s) \quad (3.8)$$

Estimation Combined with Proposition 2, Equation 3.7 suggests a very simple plug-in estimator, in which population probabilities and expectations are replaced by their empirical analogues. The same also goes for the estimation of $\bar{\tau}(S)$.

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

3.4.3 Simulation exercise

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

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

$$\Lambda(t) = \mathbb{E}[Y_{it} | \underline{T}(g) \leq t \leq \overline{T}(g)] - \sum_{s \geq 0} \mathbb{P}(t = C_i + s, C_i > 1 | \underline{T}(G_i) \leq t \leq \overline{T}(G_i)) \tau(s) \quad (3.9)$$

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

3.5 Results

3.5.1 Children-related labor supply decisions

Figure 3.4 displays the results of my event-study framework, that is average changes in hours worked by mothers whose first child is s years old, relative to (i) 2 years before this child was born, and (ii) the change in labor supplied by mothers to be over

the same time-period. Specifically, I consider (a) hours worked and (b) the salaried employment rate. Under a parallel trend and a limited anticipation assumption, these quantities identify the causal impact of motherhood on mothers' labor supply, that is the average treatment effect for a certain subpopulation (see Section 3.4).

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

Second, after they have children, the dynamics of mothers' labor supply diverge depending on the timing of their first child's birth. Specifically, after the birth of their first child, mothers decrease their hours worked relative to mothers who will have their first child later on. This decrease corresponds both to the causal effect of the extensive margin of fertility, i.e. the decision to become a mother, and to that of the intensive margin of fertility, i.e. the decision to have additional children after the first one is born. Specifically, the short-run effects ($s = 0$) correspond mostly to that of the extensive margin. As time goes by ($s > 0$), these estimates put more and more weight on the intensive margin, given that most mothers choose to have additional children, as shown in Table 3.2.

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

Figure 3.5 further decomposes the consequences of motherhood by replicating the analysis for different work-setting specific measures of labor supply. Specifically, I consider as my outcome hours worked in the salaried sector, measured in full-time units, contrasting: (a) hospital and non-hospital jobs; (b) nurse jobs and others jobs; and (c) public sector and private sector jobs. The figure makes it very clear that the impact of motherhood concentrates on public hospital nurse jobs: the effect on non-hospital, non-nurse or private sector jobs is close to 0. Part of the reason for this

is that hours worked is a non-negative quantity, and that the baseline is not high (see Figure 3.3), which mechanically constrains the magnitude of potential drops.⁷ However, the key message is that hospital nurses do not reallocate to other, plausibly more family-friendly job, upon becoming mothers: this would imply positive effects on hours worked in non-nurses, non-hospital jobs, which are rejected on the data.

In Appendix 3.C.2, I delve further into the data to expand these results. The main lessons are that (i) these aggregate results are presumably driven by the extensive margin of fertility decisions; (ii) these labor supply decisions result in substantial earnings drop; (iii) motherhood does not seem to affect overtime or working conditions, as this would lead to decreases in hourly wages that are not observed; and (iv) men do not seem to decrease their labor supply upon becoming fathers, even though the related estimates are very imprecise due to the small sample size.

3.5.2 Contribution to lifecycle profiles of labor supply

I quantify the contribution of these children-related labor supply decisions to the decline of nurses' labor supply over the course of their career by comparing the observed profiles to counterfactual ones, if female nurses did not decrease their labor supply upon becoming mothers. In practice, I simply subtract the causal effect of children, weighted by the share of mothers, from the observed profiles (see Subsection 3.4.3).

Figure 3.6 displays the results of this simulation exercise when it comes to (a) hours worked in the salaried sector, measured in full-time units, and (b) the salaried employment rate. As for hours worked, the dashed line, which represents the counterfactual lifecycle profile, progressively diverges from the plain line which corresponds to the observed profile, similar to that of Figure 3.2. However, shutting down mothers' labor supply decisions is not sufficient to rub away the decay in hours worked over the course of a career: nurses' labor supply directed towards the salaried sector would still decline even if either (i) nurses did not have children or (ii) female nurses made the same decisions as male nurses who do not reduce their hours worked upon becoming fathers. Specifically, in the simulated counterfactual,

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

average hours worked would decrease by about 0.10 full-time units over the first 10 years of a career, against an observed 0.16. In other words, children explain slightly more than a third of the decline in hours worked over the 10 first years of a career. By contrast, and consistent with Figure 3.4, the salaried employment rate of nurses would stay the same even if nurses did not have children.

Figure 3.7 further details these results by replicating this comparison across different work settings. Once again, I contrast (a) hours worked inside or outside a hospital; (b) hours worked as a nurse or not; and (c) hours worked in the public and private sector. Because the effect of children only transits through hours worked at a hospital, as a nurse or in the public sector, the counterfactual and observed profiles are the same outside these work settings. Interestingly, the magnitude of the decline would be about 0.04 full-time units in the public sector absent motherhood, against a realized 0.1: children-related labor supply decisions thus explain more than half of the decay of hours worked in the public sector.

Appendix 3.C.4 expands these results by considering the extensive margin of labor supply across work settings. Children contribute very little to the decline in nurses' salaried employment rate, except perhaps in the public sector where motherhood explains about 20% of the decline in the probability to hold a job after 10 years.

3.5.3 Robustness checks

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

Left-censoring In the DADS data, the beginning of the hospital nurse career is not observed for those who began their career in 1988 or before. This issue is made salient by the fact that the parallel trends assumption (Assumption 5) is only made conditional on both the year of birth and the timing of the beginning of the career. In other words, mothers' counterfactual labor supply if they did not have children is only imputed based on women who were born at the same time and started their career the same year as them. In practice, I choose to gather all individuals who started their career in 1988 or before in one single group, so that among them the parallel trend assumption is only conditional on the year of

birth. Figures 3.D.1 and 3.D.2 assess the robustness of my results with respect to this choice by replicating Figures 3.4 and 3.6 while omitting the data on these individuals. The results look extremely similar, confirming that they are immune to this particular issue.

Right-censoring The DADS-EDP dataset does not report childbirth that occurred after 2017. As a result, part of my control groups of mothers who are to have children later are not observed, as they cannot be differentiated from women who remain childless. This impedes the identification of cohort-specific ATTs for the year 2017, and more generally changes the composition of the control groups as one looks at younger individuals. I check that my results are not affected by this issue by replicating Figures 3.4 and 3.6 while restricting to individuals for which the distinction between mothers who have children later and childless women is less problematic. Specifically, I consider this to be the case for individuals who got their first job as a hospital nurse before 2003, as they are likely to have completed their fertility decisions by 2017, at least 15 years after they became a hospital nurse (see Figure 3.1). Figures 3.D.3 and 3.D.4 display my results. They are very close to my baseline results.

Hours worked measurement Hours worked are not observed before 1995 in the DADS data. Before this, I impute hours worked based on days worked and working-time status (full-time or part-time), under the assumption that part-time workers are on a 50% schedule. This schedule is the most frequent among part-time workers when hours are observed (see Appendix 3.A). To assess the robustness of my results with respect to this imputation, I replicate my analysis only while restricting to the 1995-2017 time-period. This will restrict the set of cohort-specific ATTs that can be identified under my identifying assumptions. Figures 3.D.5 and 3.D.6 display my results. They are once again very similar to my baseline results, which suggest that my baseline results are not driven by my choice of imputation.

Stable control group The event-study analysis upon which these results are based uses most individuals as both treated and control units: until two years before they give birth to their first child, all nurses belong to the control group, after which they become part of the treated group. As a result, the control group changes over time relative to the first child's birth: whereas, within a cohort, the short-run effect

of motherhood relies on a control group that gathers almost all nurses but the one who become mothers immediately after they get their first job, the long-run effect relies on a control group that is restricted to nurses who become mothers long after they get their first job. To assess whether the dynamics of the treatment effects is driven by this compositional shift in the control group, I replicate my analysis, this time restricting the control group to nurses who have their first child at least 14 years after they got their first nurse job at a hospital, and restricting the treated group to nurses who have their first child at most 13 years after their first job. Additionally, I restrict the data to the 13 years that follow the first hospital nurse job. This approach is akin to more traditional difference-in-difference approaches in which units do not switch from one group to the other over time. Figure 3.D.7 displays the resulting event-study estimates. They are very close to my baseline estimates, especially regarding hours worked. The employment effects appear slightly larger but are still compatible with my baseline results. Furthermore, Figure 3.D.8 shows that even so, the contribution of motherhood to the labor supply lifecycle at the extensive margin remains marginal at best.

3.6 Conclusion

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

These findings imply that the supply of nursing labor should be studied in relation to fertility decisions, gender norms and family policies, as suggested by [Friedrich and Hackmann \(2021\)](#). Specifically, they suggest that the aggregate supply of nurs-

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


References

- Abadie, A., Athey, S., Imbens, G. W., and Wooldridge, J. (2017). When should you adjust standard errors for clustering? Working Paper 24003, National Bureau of Economic Research.
- Antonazzo, E., Scott, A., Skatun, D., and Elliott, R. F. (2003). The labour market for nursing: a review of the labour supply literature. *Health Economics*, 12(6):465–478.
- Askildsen, J. E., Baltagi, B. H., and Holmås, T. H. (2003). Wage policy in the health care sector: a panel data analysis of nurses' labour supply. *Health Economics*, 12(9):705–719.
- Barlet, M. and Cavillon, M. (2010). La profession infirmière: situation démographique et trajectoires professionnelles. Études et Résultats 759, Drees.
- Bessière, S. (2005). La féminisation des professions de santé en france: données de cadrage. *Revue française des affaires sociales*, 1:17–33.
- Callaway, B. and SantAnna, P. H. (2020). Difference-in-differences with multiple time periods. *Journal of Econometrics*.
- Daune-Richard, A.-M., Odena, S., and Petrella, F. (2007). Entreprises et modes d'accueil de la petite enfance. innovation et diversification. Dossier d'étude 91, Cnaf.
- de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–96.
- Estryn-Béhar, M., der Heijden, B. I. J. M. V., Ogiska, H., Camerino, D., Nézet, O. L., Conway, P. M., Fry, C., and Hasselhorn, H.-M. (2007). The impact of social work environment, teamwork characteristics, burnout, and personal factors upon intent to leave among european nurses. *Medical Care*, 45(10):939–950.
- Friedrich, B. U. and Hackmann, M. B. (2021). The Returns to Nursing: Evidence from a Parental-Leave Program. *The Review of Economic Studies*.
- Frijters, P., Shields, M. A., and Price, S. W. (2007). Investigating the quitting decision of nurses: panel data evidence from the british national health service. *Health Economics*, 16(1):57–73.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*.
- Gruber, J. and Kleiner, S. A. (2012). Do strikes kill? evidence from new york state. *American Economic Journal: Economic Policy*, 4(1):127–57.
- Hanel, B., Kalb, G., and Scott, A. (2014). Nurses labour supply elasticities: The importance of accounting for extensive margins. *Journal of Health Economics*, 33(C):94–112.
- Holmås, T. H. (2002). Keeping nurses at work: a duration analysis. *Health Economics*, 11(6):493–503.
- Kleven, H., Landais, C., and Sogaard, J. E. (2019). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*,

- 11(4):181–209.
- Lafortune, G., Socha-Dietrich, K., and Vickstrom, E. (2019). Recent trends in international mobility of doctors and nurses. In OECD, editor, *Recent Trends in International Migration of Doctors, Nurses and Medical Students*, pages 11–34. OECD Publishing.
- Nooney, J. G., Unruh, L., and Yore, M. M. (2010). Should I stay or should I go? Career change and labor force separation among registered nurses in the U.S. *Social Science & Medicine*, 70(12):1874–1881.
- Phillips, V. L. (1995). Nurses’ labor supply: Participation, hours of work, and discontinuities in the supply function. *Journal of Health Economics*, 14(5):567–582.
- Pora, P. and Wilner, L. (2019). Child penalties and financial incentives: Exploiting variation along the wage distribution. Documents de Travail de l’Insee - INSEE Working Papers G2019/08, Insee.
- Propper, C. and VanăReenen, J. (2010). Can pay regulation kill? panel data evidence on the effect of labor markets on hospital performance. *Journal of Political Economy*, 118(2):222–273.
- Reynaud, D. (2020). Les femmes les plus modestes et les plus aisées ont le plus denfants. Insee Première 1826, Insee.
- Shields, M. A. (2004). Addressing nurse shortages: what can policy makers learn from the econometric evidence on nurse labour supply?*. *The Economic Journal*, 114(499):F464–F498.
- Shields, M. A. and Ward, M. (2001). Improving nurse retention in the National Health Service in England: the impact of job satisfaction on intentions to quit. *Journal of Health Economics*, 20(5):677–701.
- Sun, L. and Abraham, S. (2020). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*.
- Toren, O., Zelker, R., Lipschuetz, M., Riba, S., Reicher, S., and Nirel, N. (2012). Turnover of registered nurses in Israel: Characteristics and predictors. *Health Policy*, 105(2):203–213.
- Wilner, L. (2016). Worker-firm matching and the parenthood pay gap: Evidence from linked employer-employee data. *Journal of Population Economics*, 29(4):991–1023.

Figures

Figure 3.1 – Lifecycle profile of fertility: share of mothers among nurses of both genders

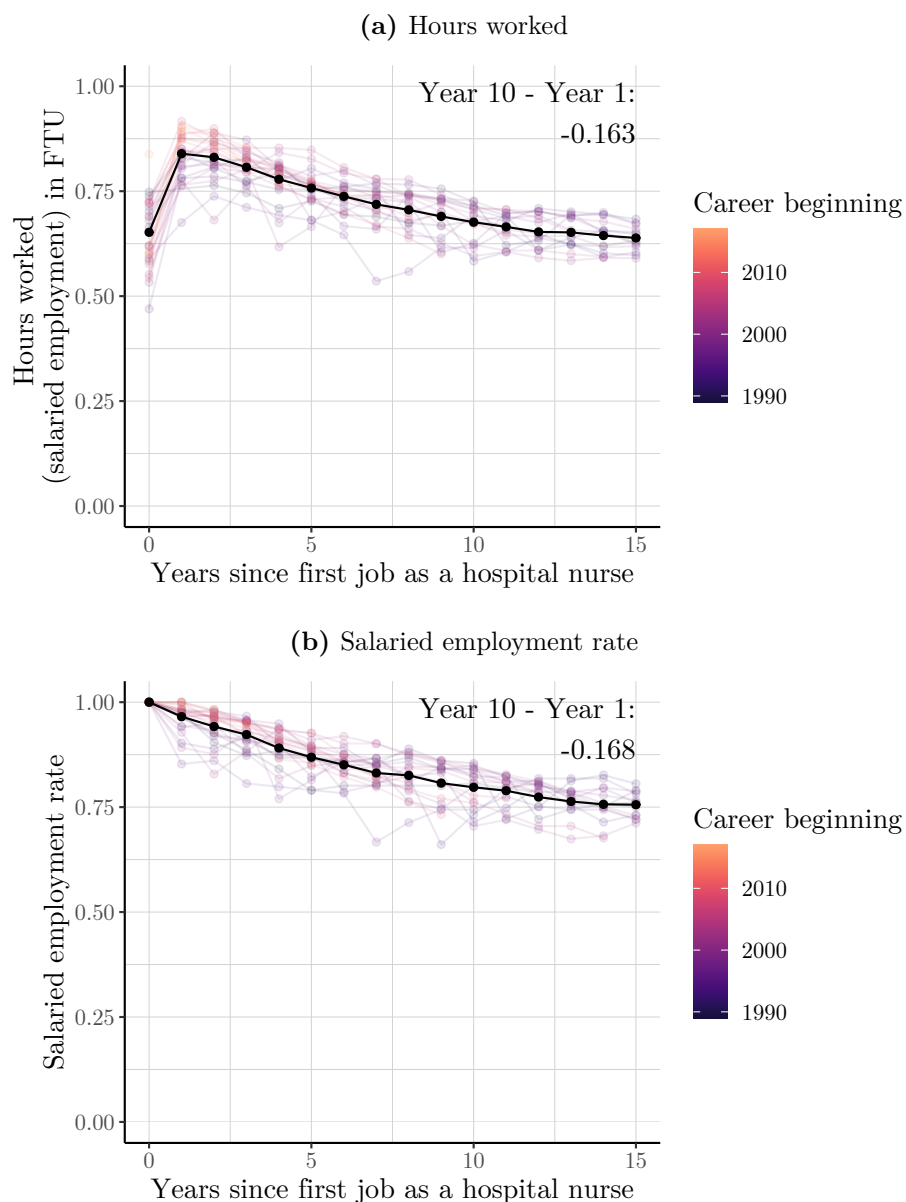


Share of mothers, by time relative to the first qualified healthcare worker job at a hospital.

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

Source. Insee, DADS-EDP panel.

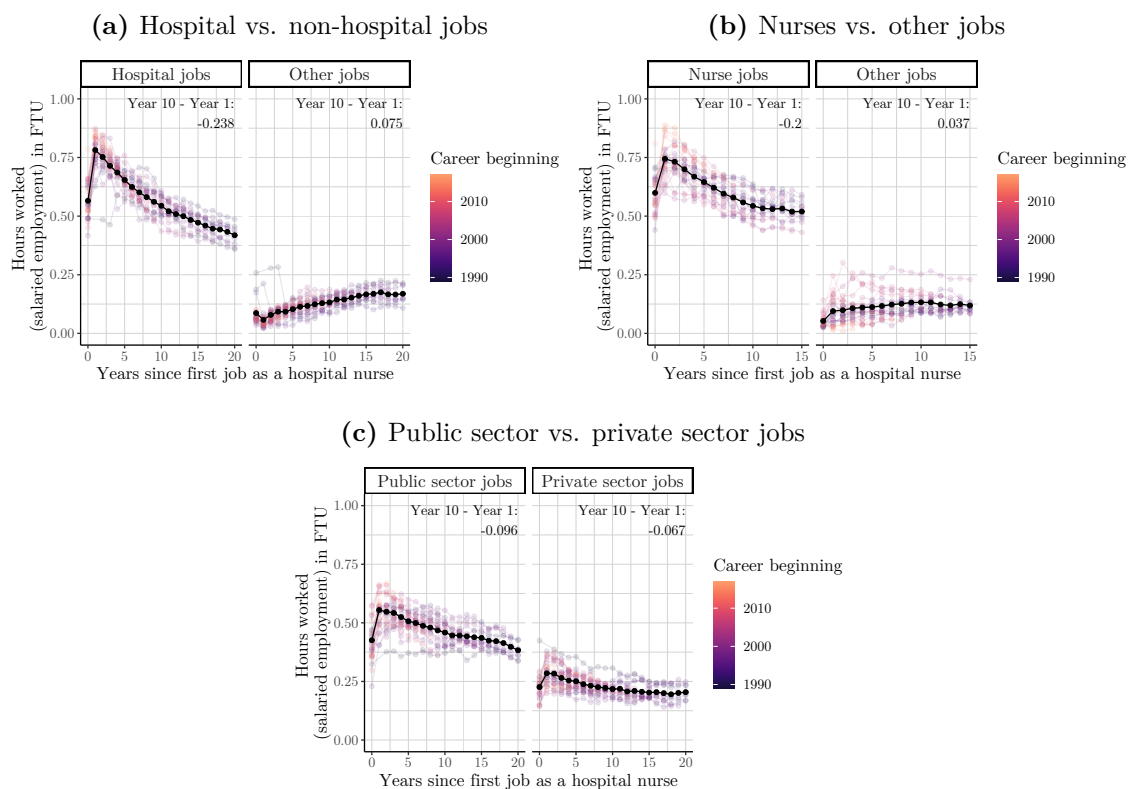
Figure 3.2 – Lifecycle profile of hospital nurses' labor supply: total labor supply in the salaried sector



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

Source. Insee, DADS panel.

Figure 3.3 – Lifecycle profile of hospital nurses' labor supply: decompositions




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

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

Source. Insee, DADS-EDP panel.

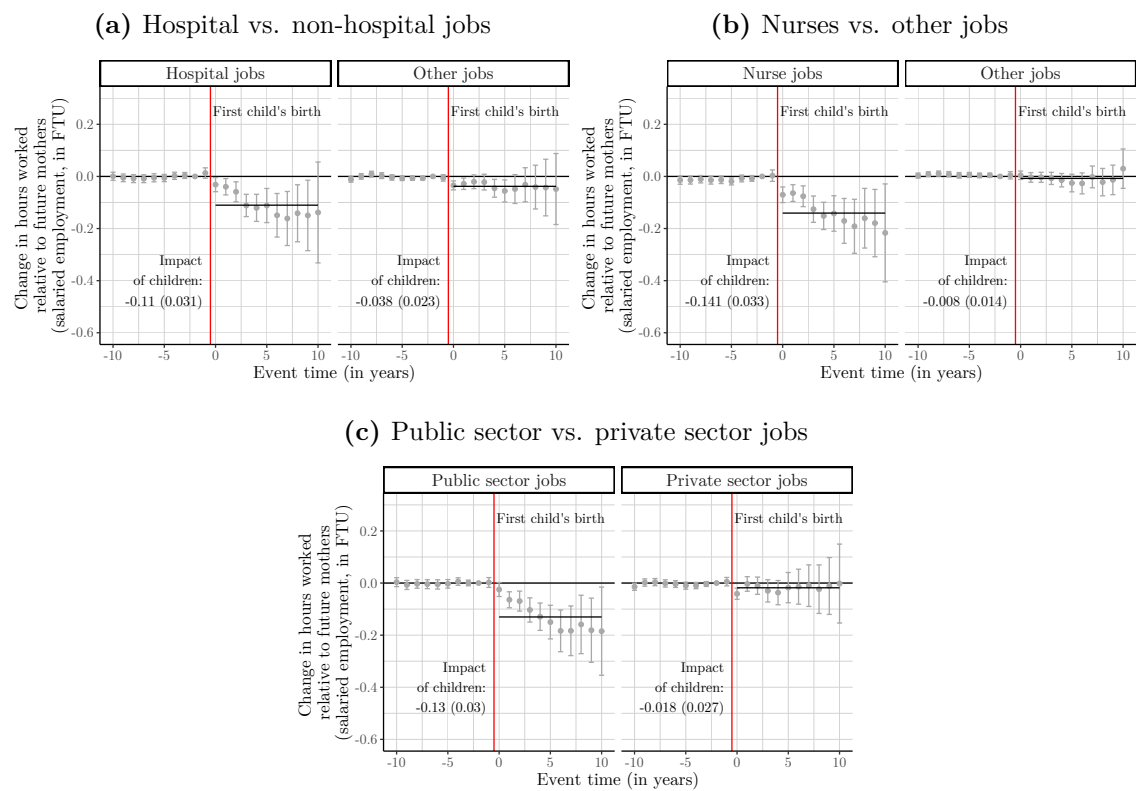
Figure 3.4 – Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector



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

Source. Insee, DADS-EDP panel.


Figure 3.5 – Event-study estimates of the impact of children on mothers’ labor supply: decompositions



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

Source. Insee, DADS-EDP panel.

Figure 3.6 – Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector

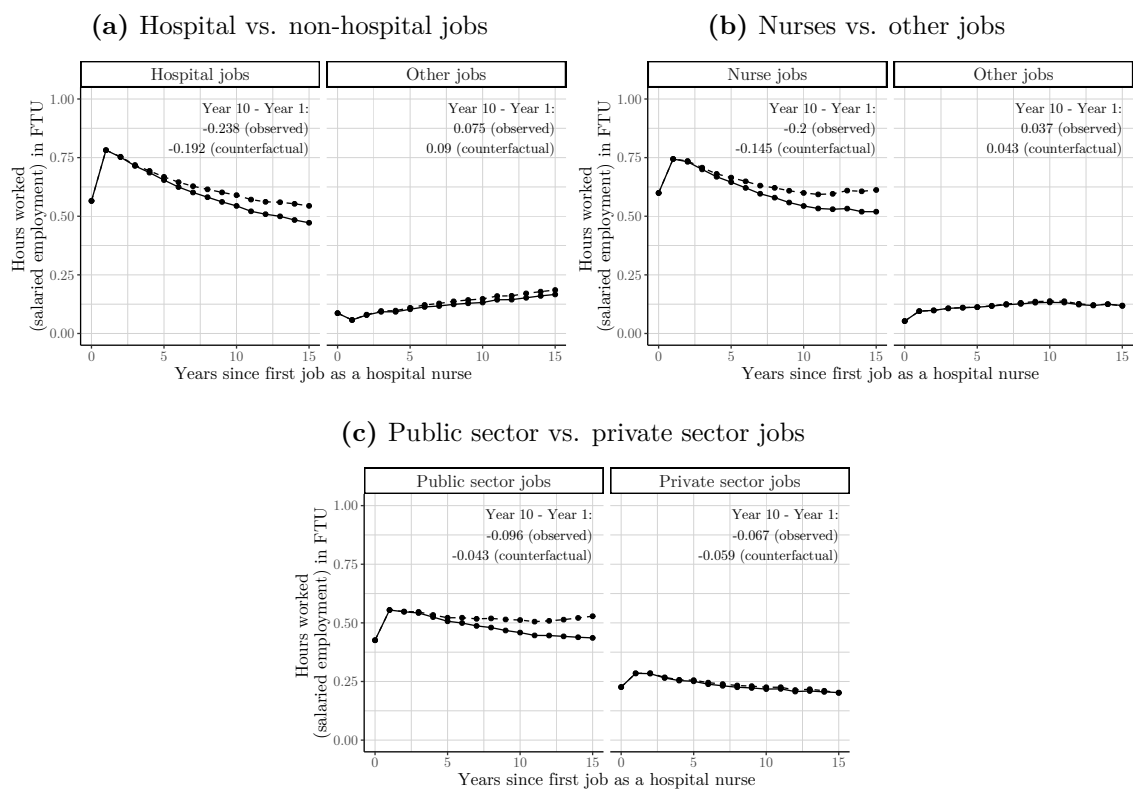


Realized and counterfactual average hours worked in the salaried sector, in full-time units, and salaried employment rate, by time relative to the first qualified healthcare worker job at a hospital. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0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

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

Source. Insee, DADS-EDP panel.

Figure 3.7 – Contribution of children to the lifecycle profile of nurses' labor supply: decompositions



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

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

Source. Insee, DADS-EDP panel.

Tables

Table 3.1 – Detailed occupations distribution among selected jobs (2009-2017)

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

Source. DADS panel, Insee.

Table 3.2 – Summary statistics

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

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

3.A DADS panel: labor supply measures

3.A.1 Hours worked: concept

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

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

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

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

3.A.2 Full-time units conversion

Hours worked are converted in full-time units using a very simple approach, that relies on three variables: working-time status, days worked and lastly hours worked. This approach caps time worked at 1 for individuals who work full-time during an

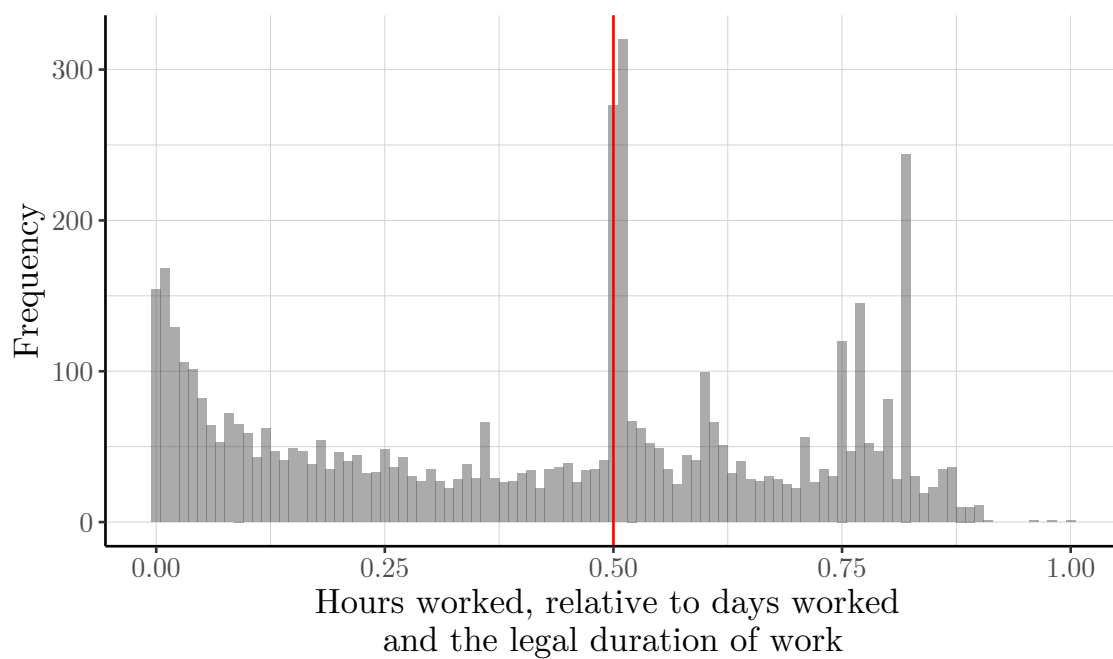
entire year, so that it does not incorporate overtime. The main advantage of this method is that it allows to compare time worked even when the legal duration of work changes, which is the case over my time period of interest as this duration changed from 39h to 35h per week in the beginning of the 2000s.

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

An issue with this approach is that hours worked are not observed before 1995 in the DADS data. Before this date, only the working-time status and days worked can be observed. I choose to impute full-time units to part-time workers under the assumption that they are on a 50% schedule, so that their time worked is proportional to days worked and equal to 0.5 if they work an entire year. I make this choice because 50% schedule was likely the most frequent choice among part-time workers. Figure 3.A.1 makes it salient by plotting the distribution of hours worked, relative to (i) the legal duration of work and (ii) days worked, among part-time job spells observed between 1995 and 1998 for workers who at one point in time were observed in hospital nurses positions.

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

Figure 3.A.1 – Distribution of hours worked for part-time workers, 1995-1998



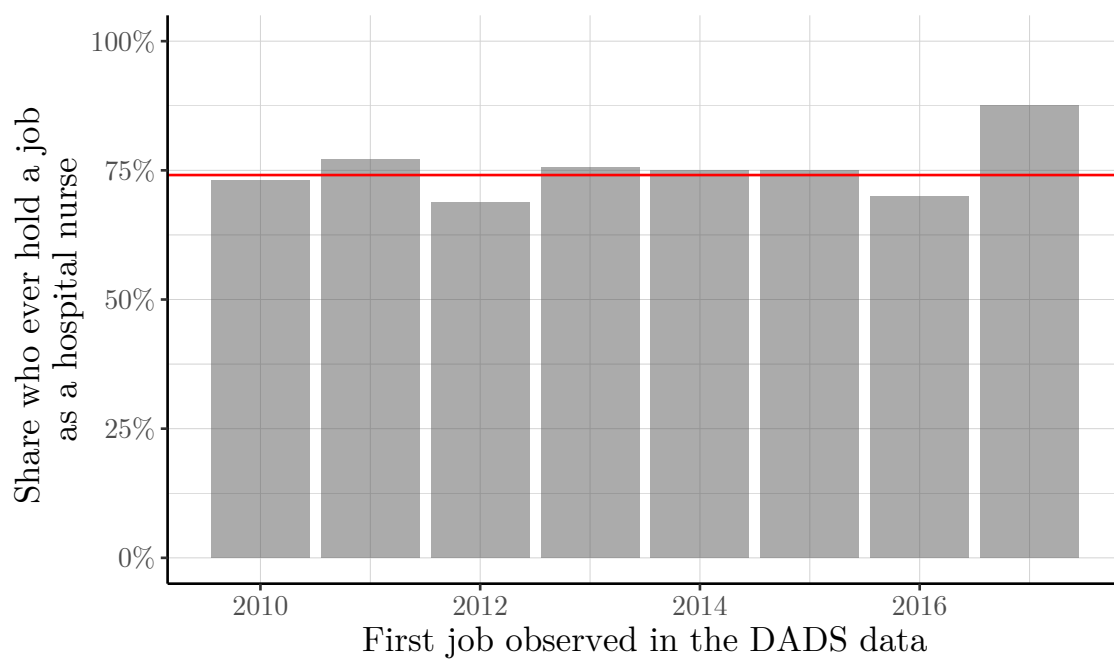
Distribution of hours worked divided by the legal annual duration of work for full-time workers, among individuals who hold a hospital nurse job at least once.

Source. Insee, DADS panel.

3.B How well does the sample cover the nurse occupation?

Table 3.1 shows that even though I cannot rely on the detailed occupation variable to delineate my sample, it mostly covers individuals that can reasonably be considered as hospital nurses. Reciprocally, a question is how many nurses do, at some point of their lives, hold a job as a hospital nurse, and thus fall within the universe that I intend to cover. To investigate this issue, I consider all individuals (i) observed as holding a salaried nurse job, in terms of the detailed occupation variable, at any point in time since 2010; and (ii) whose first job, whatever the industry or occupation, started in 2010 or later. The second condition is meant to capture individuals whose entire career can be observed in terms of the detailed occupation variable, so as to avoid selection issue that would stem from nurses gradually leaving the occupation over time (see Figure 3.3). Because nurses cannot begin their careers as freelance nurses, but have to hold a job as a salaried nurses, these conditions should include all individuals who began their careers in 2010 or after and hold a nurse job at some point. I then compute the share of these individuals who fall within my sample. This measure gives a lower bound of the share of nurses who hold a hospital nurse job at some point of their lives, given that I can only observe the very beginning of a career for the selected individuals.

Figure 3.B.1 displays my estimates. Even though there is variation from one cohort to the other, which is probably due to small sample size, it suggests that at least 74% of nurses hold a job as a nurse at a hospital at some point of their lives.

Figure 3.B.1 – Share of nurses who ever hold a job at a hospital


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

Source. Insee, DADS panel.

3.C Additional results

3.C.1 Lifecycle profiles

Figure 3.C.1 – Lifecycle profile of hospital nurses' labor supply: decompositions at the extensive margin



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

Source. Insee, DADS-EDP panel.

3.C.2 Children-related labor supply decisions

Extensive and intensive margins of fertility decisions

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

To gain further insights into this issue, I replicate my event-study analysis (a) only relying on observations related to mothers of one or mothers of additional children *at least two years before their second child is born*; and (b) considering the second child's birth instead of the first child's birth as my event of interest *and* only relying on observations related to mothers of two or mothers of additional children *at least two years before their third child is born*. Under parallel trends and limited anticipation assumptions very similar to Assumptions 5 and 6,⁹ this allows to distinguish between (a) the dynamic effect of the first child and (b) the dynamic effect of having one additional child (i.e. the second child), i.e. the dynamic effect of (a) the extensive and (b) the intensive margin of fertility, without contamination from subsequent fertility decisions.

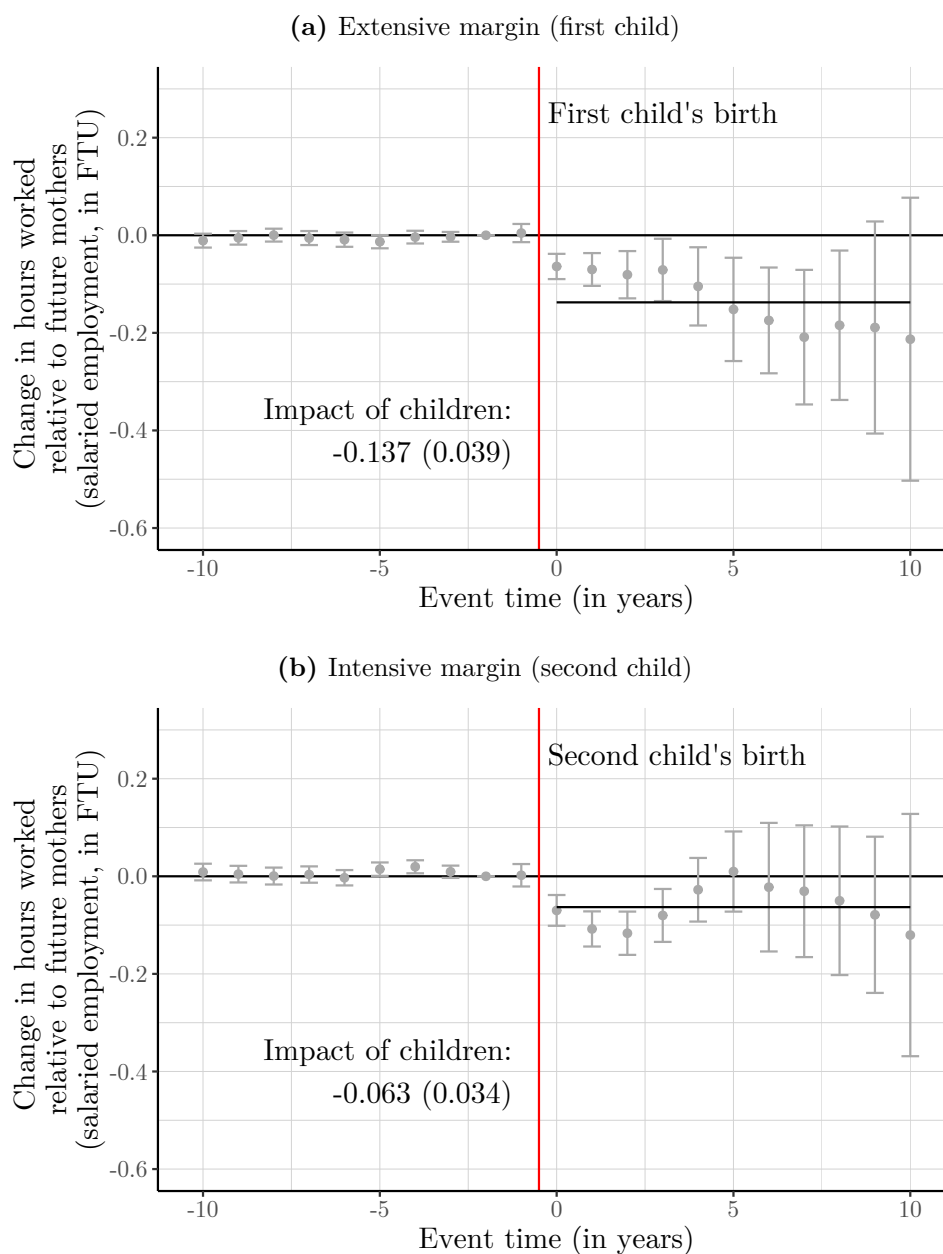
Figure 3.C.2 displays these estimates. While they may be slightly less precise than the baseline estimates, because they rely on less observations, they are still informative to some extent. Specifically, they show that while both margins seem to

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

have somewhat similar short-run effects, the impact of having one additional child is quite short-lived. In other words, the impact of the second child seems to vanish after a few years, whereas the impact of the first one is long-lasting. Moreover, the magnitude of the impact of the first child is not very different from that of my baseline estimates that mix up all margins.

This suggests that these baseline results are mostly driven by the long-run consequences of fertility decisions at the extensive margin, i.e. the decisions to become a mother, rather than the short-run consequences of fertility decisions at the intensive margin, i.e. the decision to have one additional child among women who are already with child. These results are also, to some extent, informative about the origins of such labor supply decisions. Indeed, if these decisions were purely the result of children-related time constraint, as opposed to norms and preferences, one would expect the dynamic path of the effect of the first child to be very similar to that of the second child, because conditional on their age, the needs of children should not be strongly dependent on their rank among their siblings. That they do really differ therefore seems to indicate either (i) large returns to scale in the children production function or (ii) that motherhood-related time allocation decisions involve gender norms and preferences.

Figure 3.C.2 – Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector, by fertility margin

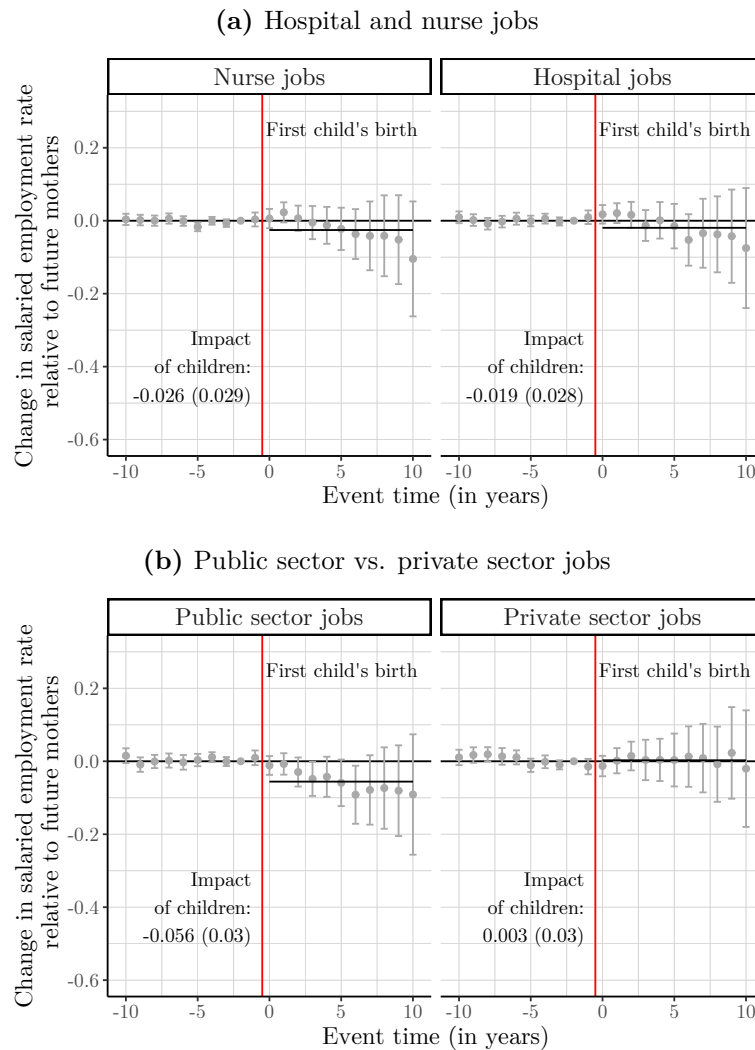


Event-study estimates of the impact of children on mothers' hours worked in the salaried sector, in full-time units, by time since first child's birth. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0 hours worked). Standard errors are clustered at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

Source. Insee, DADS-EDP panel.

Other margins of labor supply


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



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

Source. Insee, DADS-EDP panel.

Figure 3.C.4 – Event-study estimates of the impact of children on mothers' labor outcomes: labor earnings

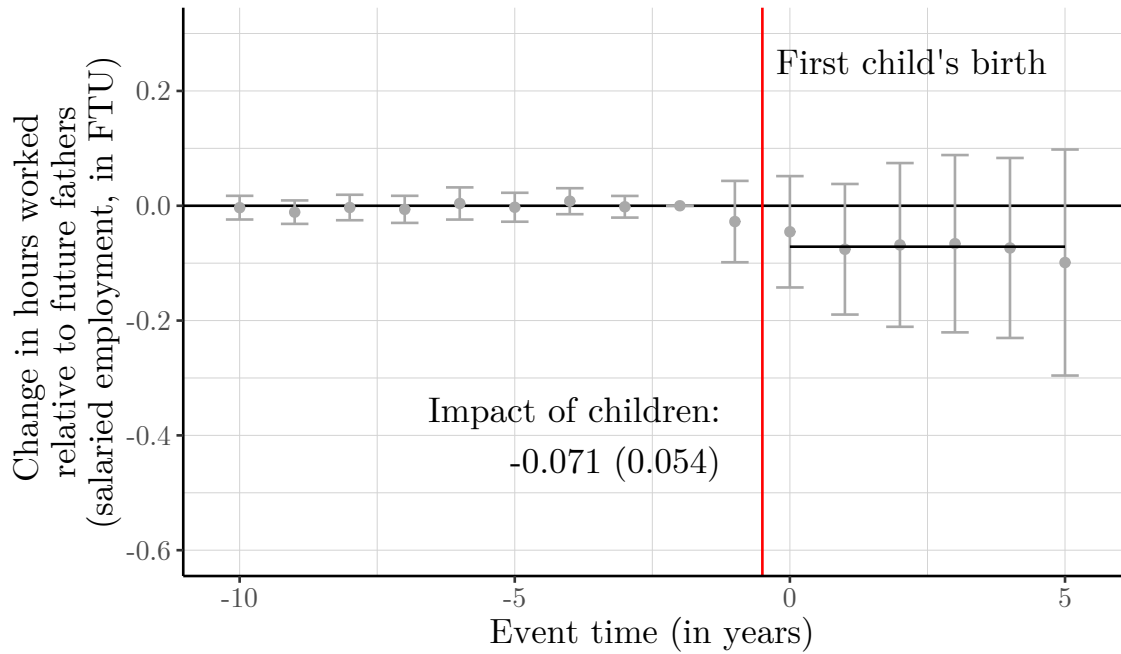


Event-study estimates of the impact of children on mothers' salaried labor earnings, by time since first child's birth. Standard errors are clustered at the individual level and estimated by bootstrap with 200 replications.

Source. Insee, DADS-EDP panel.

Children-related labor supply decisions: fathers

Figure 3.C.5 – Event-study estimates of the impact of children on fathers' labor supply: total labor supply in the salaried sector



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

Source. Insee, DADS-EDP panel.

3.C.3 Overtime, working conditions and hourly wages

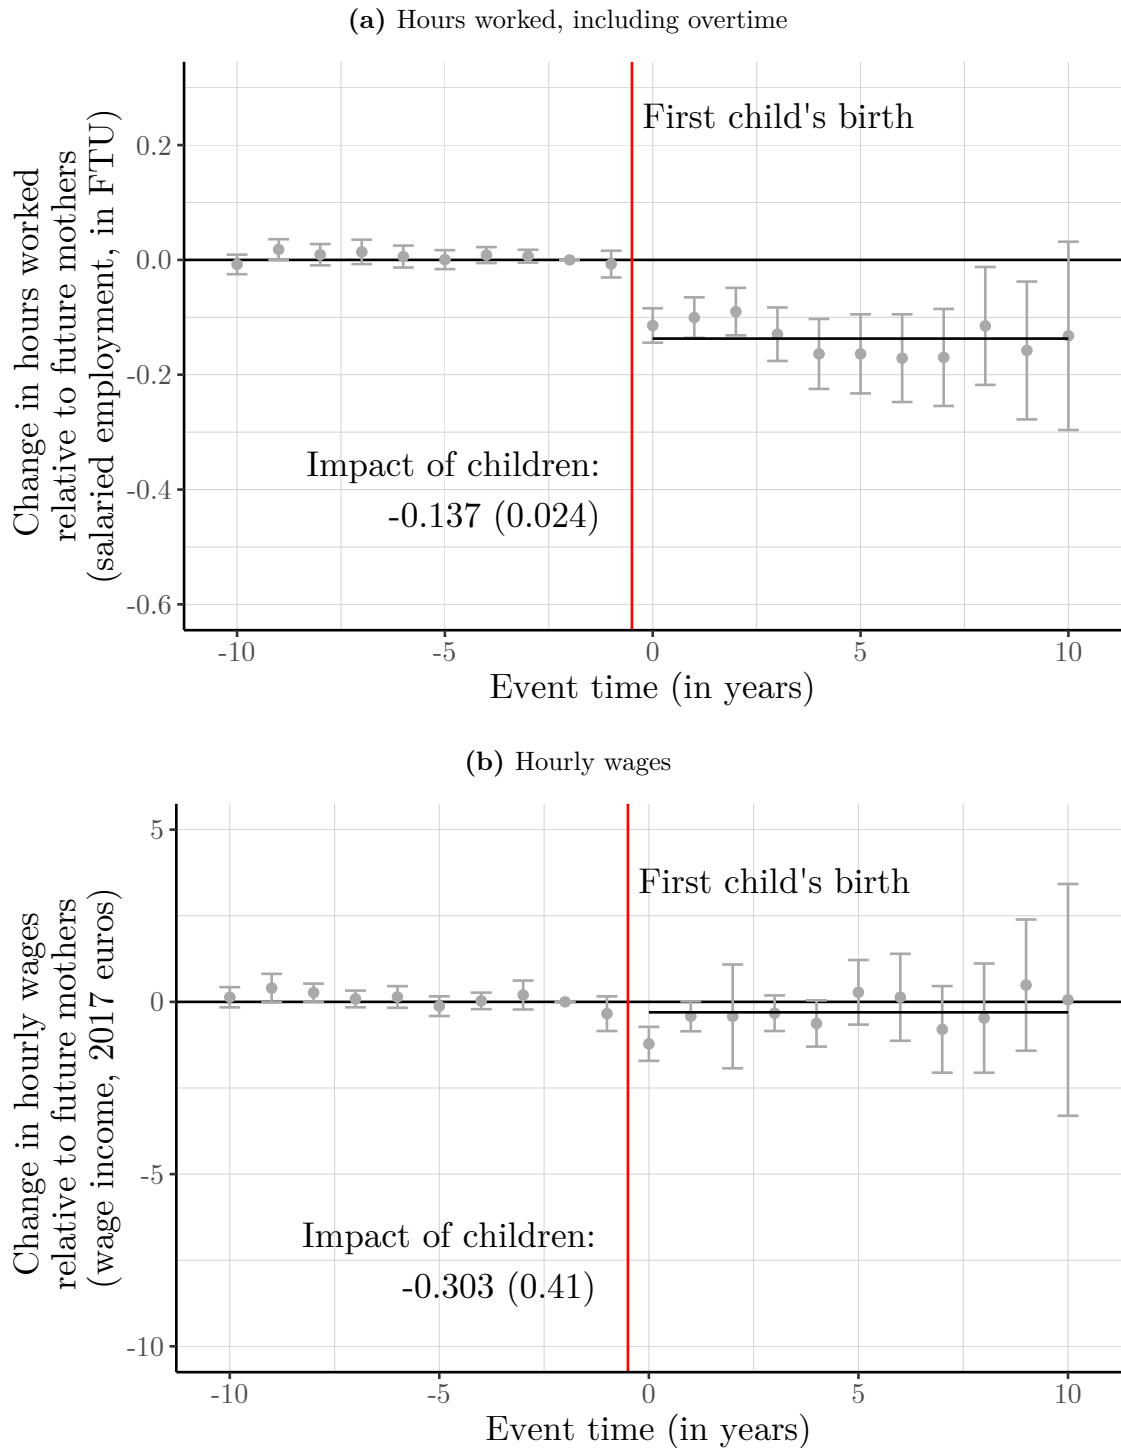
Hours worked, as measured in full-time units, does not include overtime because this measure is capped to 1 for full-time workers. As such, if some nurses choose to keep their full-time jobs, but cut their overtime hours upon becoming mothers, then my estimates of the magnitude of children-related labor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

To investigate these dimensions, in Figure 3.C.6 I replicate my event-study analysis, focusing on two new outcomes. First, I consider hours worked including overtime, relative to an entire year of full-time work.¹⁰ This measure is akin to hours worked measured in full-time units, except that it is no longer capped to 1 for full-time workers. My results are extremely similar to those obtained with hours worked measured in full-time units, which indicates that changes in overtime are not an important margin for children-related labor supply decisions.

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

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

Figure 3.C.6 – Event-study estimates of the impact of children on mothers’ labor outcomes: hours worked including overtime and hourly wages




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

Source. Insee, DADS-EDP panel.

3.C.4 Contribution to lifecycle profiles of labor supply

Figure 3.C.7 – Contribution of children to the lifecycle profile of nurses' labor supply: decompositions at the extensive margin



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

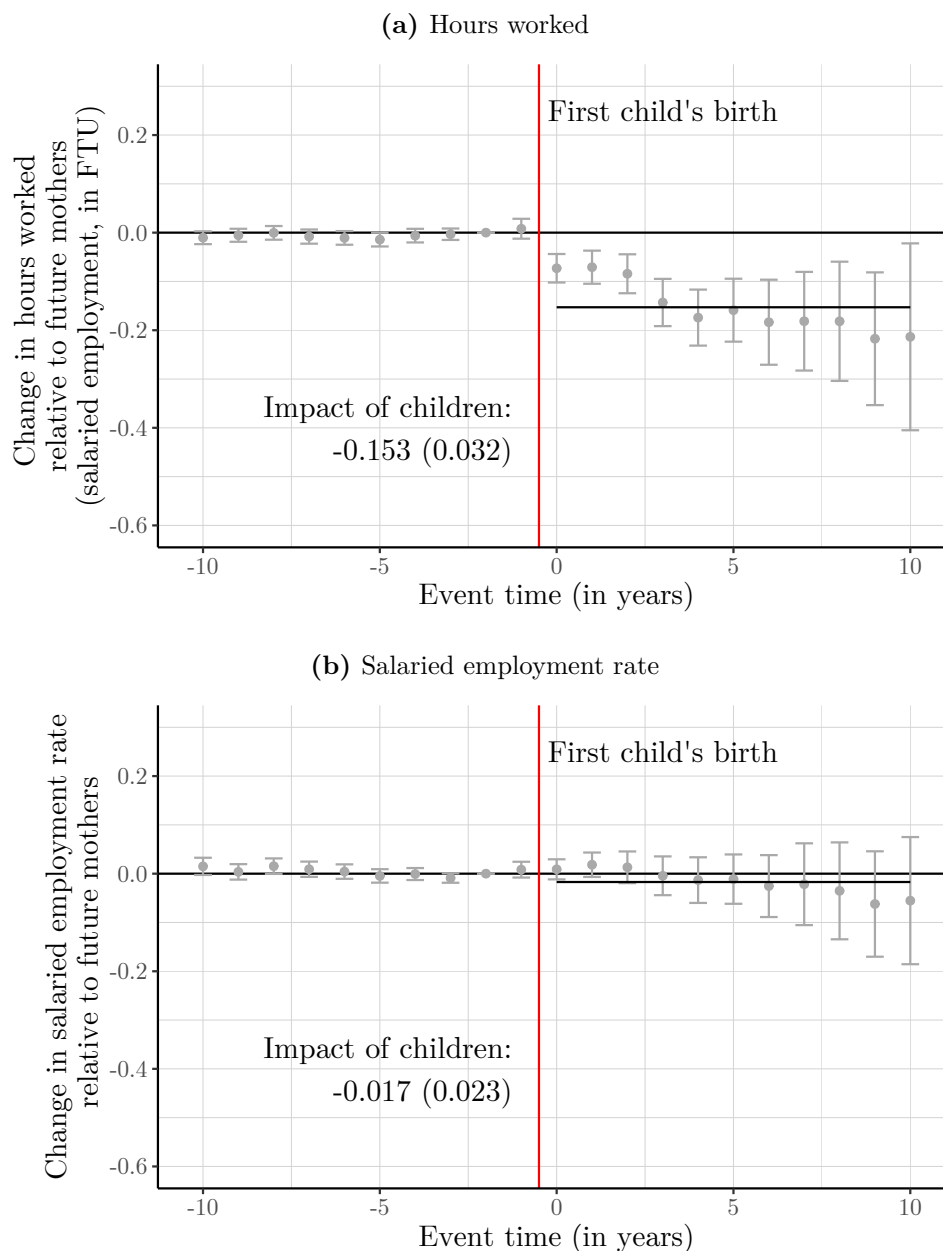
Note. Data on individuals who got their first hospital nurse job in 1988 or before are omitted from the computation.

Source. Insee, DADS-EDP panel.

3.D Robustness checks

3.D.1 Left-censoring issue


Figure 3.D.1 – Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector



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

Source. Insee, DADS-EDP panel.

Figure 3.D.2 – Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector



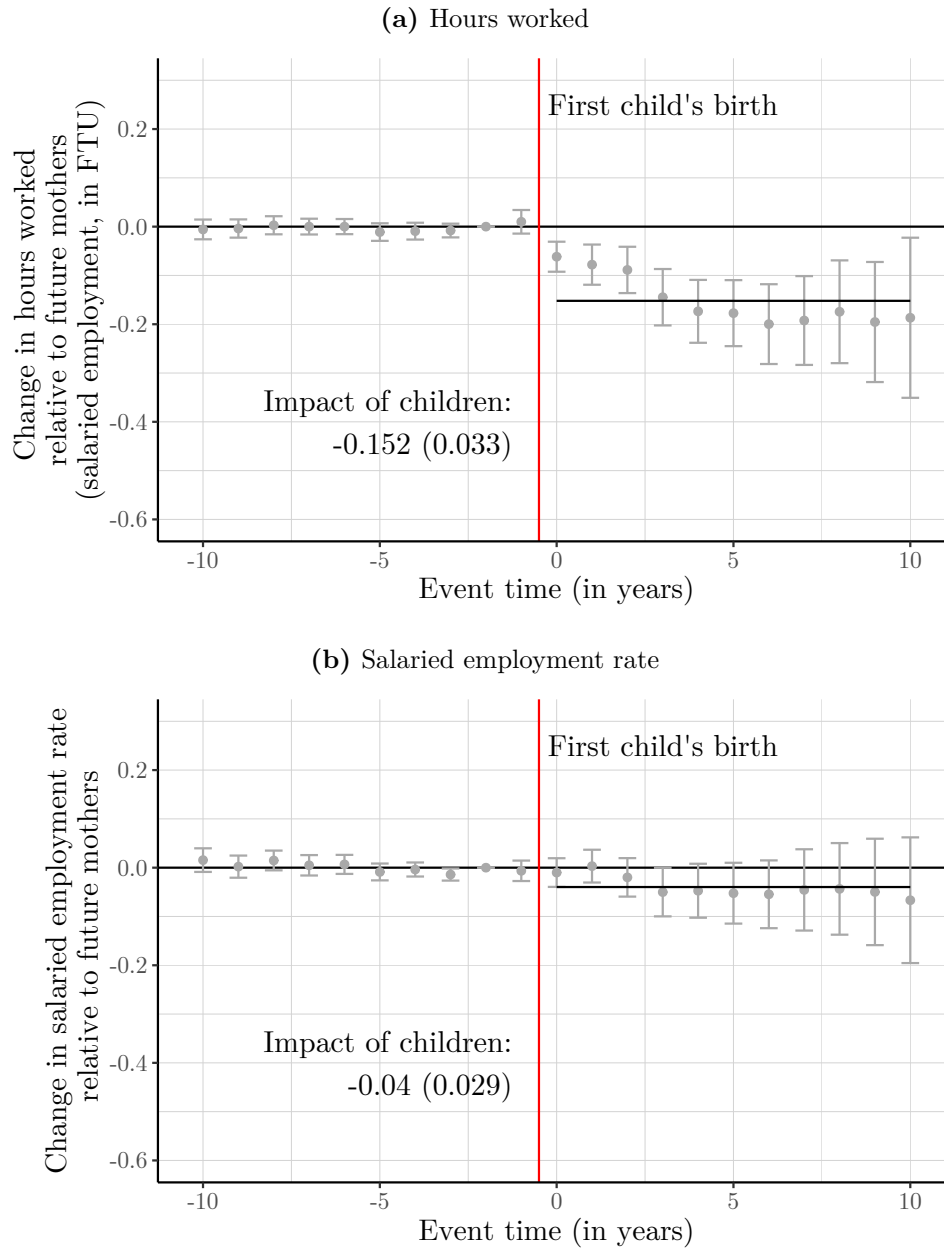
Realized and counterfactual average hours worked in the salaried sector, in full-time units, and salaried employment rate, by time relative to the first qualified healthcare worker job at a hospital. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0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

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

Source. Insee, DADS-EDP panel.

3.D.2 Right-censoring issue

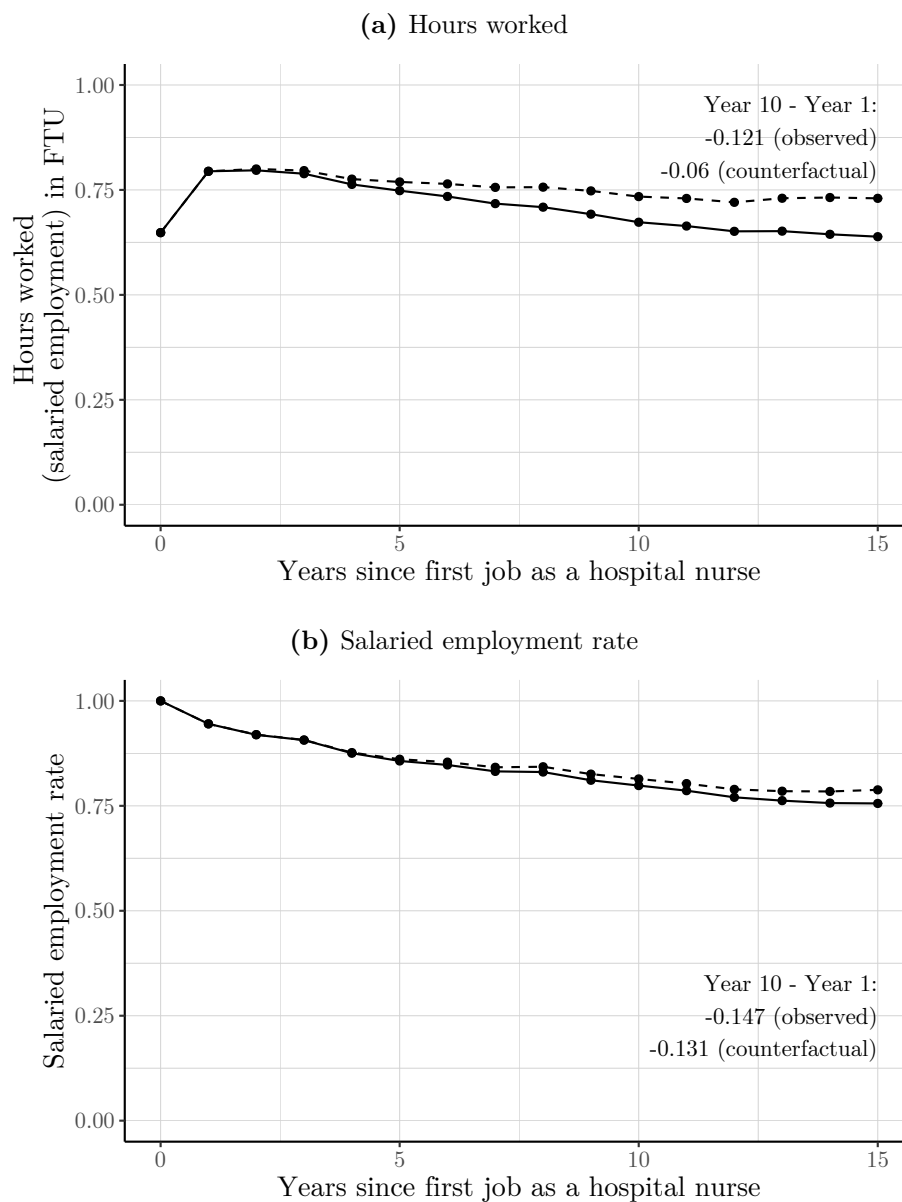
Figure 3.D.3 – Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector



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

Source. Insee, DADS-EDP panel.

Figure 3.D.4 – Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector



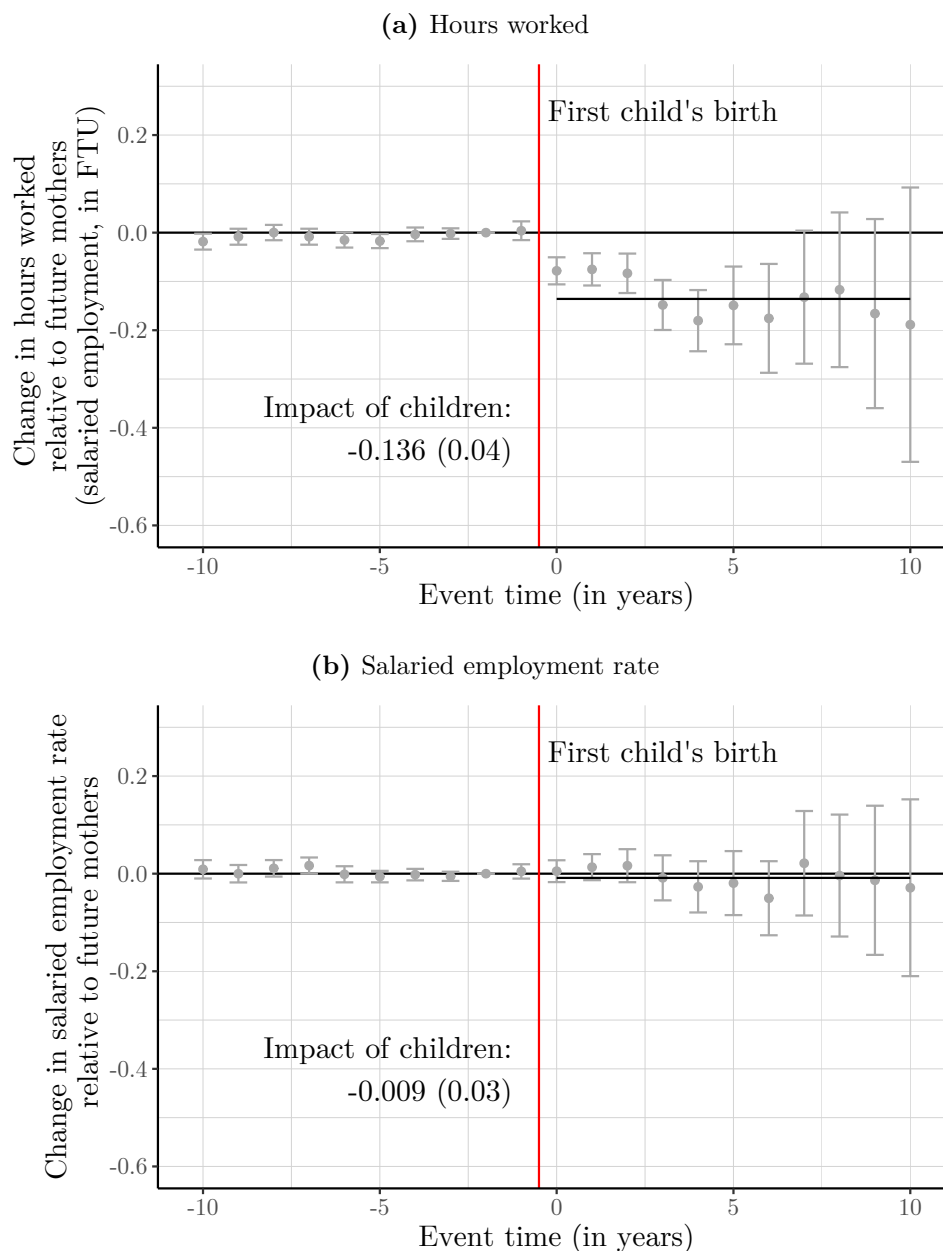
Realized and counterfactual average hours worked in the salaried sector, in full-time units, and salaried employment rate, by time relative to the first qualified healthcare worker job at a hospital. Hours worked are not conditional on salaried employment, but incorporate the participation margin (0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

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

Source. Insee, DADS-EDP panel.

3.D.3 Hours worked measurement


Figure 3.D.5 – Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector



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

Source. Insee, DADS-EDP panel.

Figure 3.D.6 – Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector



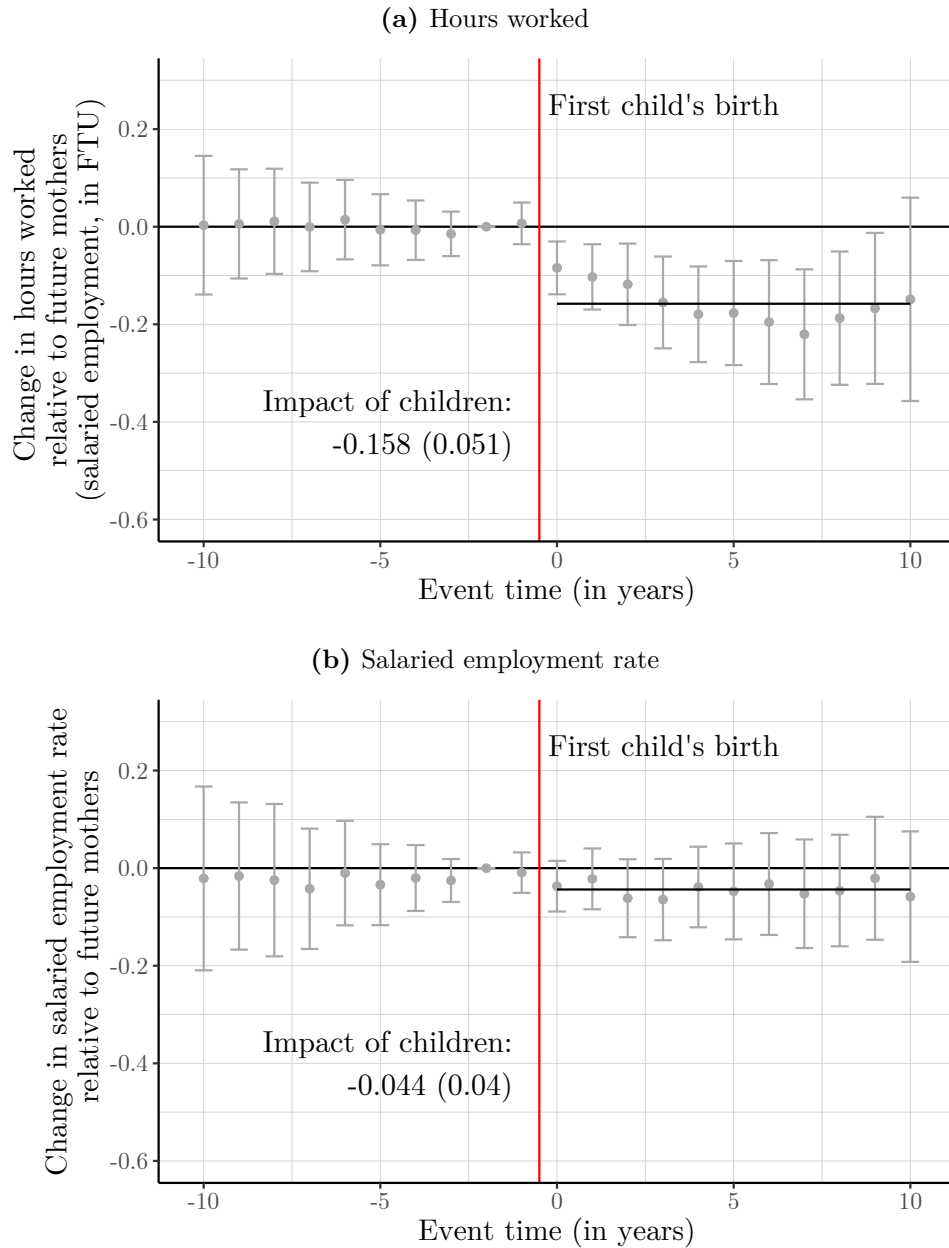
Realized and counterfactual average hours worked in the salaried sector and, in full-time units, and salaried employment rate by time relative to the first qualified healthcare worker job at a hospital. Hours worked are not conditional on salaried employment, but incorporate the participation margin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

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

Source. Insee, DADS-EDP panel.

3.D.4 Stable control group


Figure 3.D.7 – Event-study estimates of the impact of children on mothers' labor supply: total labor supply in the salaried sector



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

Source. Insee, DADS-EDP panel.

Figure 3.D.8 – Contribution of children to the lifecycle profile of nurses' labor supply: total labor supply in the salaried sector



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

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

Source. Insee, DADS-EDP panel.

Chapter 4

Job displacement, families and redistribution

joint with Raphaël Lardeux

4.1 Introduction

The reallocation of workers across jobs, especially due to innovation-related productivity shocks, has been emphasized as a necessary condition for economic growth (Aghion and Howitt, 1994). However, the collective benefits of the reallocation may be deemed unacceptable if this reallocation results in a costly externality for the restricted subset of workers who are actually hit by job displacement. This concern, either conscious or not, serves as the basis for the various safety nets and social insurance systems that characterize modern economies.

On top of these collectively organized insurance systems, workers may also mitigate the related income risk at a more micro level, thanks to their own families. Indeed, as long as the within family correlation in a earnings risk remains small, families may offer good insurance against individual idiosyncratic risk (i) by diluting the individual risk thanks to the pooling of resources, and thus of all household members' individual shocks, and (ii) by limiting the welfare implications of the said risk through the within-family reallocation of time and effort based on the new structure of within-family comparative advantage (Becker, 1981) that emerges as a consequence of a shock on the return on time spent on the labor market by one family member. This mechanism is generally coined as the added worker effect (Lundberg, 1985), and the optimal taxation of couples depends on whether or not this channel is active (Kleven et al., 2009).

Even though job displacement emphasizes the insurance role played by families, its effect on family structure, especially in the context of heterosexual couples, remains theoretically ambiguous. Indeed, while for the displaced worker, due to the family insurance mechanisms, the value of belonging to a family increases due to job displacement, the opposite is true for the non-displaced partner. Moreover, these adverse effects may be further enhanced if couples incur a psychological cost when the new allocation of time and effort is perceived to conflict traditional gender norms (Akerlof and Kranton, 2000). Given the usual gender roles in Western countries, this might particularly be the case when job displacement hit men, as it may affect their status as breadwinners.

In this paper, we investigate this issue in the French context, taking advantage of a combination of administrative records that span over multiple dimensions of family structure, individual and household income, and labor market trajectories. Building on the methodology inaugurated by Jacobson et al. (1993) and extended

by [Halla et al. \(2020\)](#), we elaborate a sample of displaced workers which are affected by a plausibly exogenous job separation, and compare them to very similar workers who are not affected by such a shock to investigate (i) whether job displacement disrupts couples and families, or decreases the odd of single workers finding a new partner; (ii) whether spouses provide insurance against the related income risk by increasing their labor supply in response to job displacement; and (iii) how well families, safety nets and capital markets insure workers against the income risk that results from job displacement.

We find that job displacement has close to no effect on average on family structure: the probability of living in a cohabiting heterosexual relationship does not change as a consequence of job separation, and workers' fertility decisions do not seem to react to the shock. We do not find evidence of an added worker effect: displaced workers' spouses' earnings and labor supply do not increase in response to the sudden decrease in their partners' earnings. In other words, if families provide insurance against the income risk that results from job displacement, it is only through the pooling of resources rather than from collective adjustments in the allocation of time and effort. Lastly, whereas unemployment insurance seems to insure workers quite effectively against the short-lived employment shock, it matters much less in the long run when most workers' have managed to find a new job. Workers seem partially insured at best against the long-run income loss, through a bundle that involves self-employment, capital markets and the progressivity of the taxes and transfers system. Even so, the magnitude of the long-run shock even after these resources are taken into account remains sizable.

These results contribute to three related strands of the literature. First, several papers have investigated the impact of job displacement on family structure, for instance in terms of the divorce probability ([Charles and Stephens, 2004](#); [Eliason, 2012](#)) or in terms of fertility decisions ([Huttunen and Kellokumpu, 2016](#)). With respect to the evidence on divorce, our investigation is slightly broader because we focus both on the probability of maintaining a cohabiting heterosexual relationship after job displacement for already formed couples and on the probability of finding a mate for single workers, i.e. both on couple dissolution and couple formation. Consistent with the American evidence, but in contrast with the Swedish case, we do not find that job displacement induced by downsizing and plant closure affects this probability. This evidence also diverges from multi-country investigations that suggest that at the individual level, male job loss correlates with higher rates of

divorce (Solaz et al., 2020), and recent results on the fall of the marriage value of young men as a result of their falling relative earnings, which are identified from more macro shocks (Autor et al., 2019a).

We also find close to no effect of job displacement on fertility decisions, with the exception of a slight decrease in men's fertility, which contrasts the Finnish evidence. Because they are informative about the effect of income shocks on fertility decisions conditional on workers' time-constant heterogeneity, these null effects also speak to methodological issues regarding the effect of children on labor supply and labor earnings. Specifically, they suggest that differences in the timing of children's birth are unlikely to stem from unobserved shocks on parents' earnings, so that event-study and difference-in-difference approach yield credible results (Kleven et al., 2019).

Second, we contribute to the literature that has investigated within-family time and effort reallocation in response to individual income shock, which has long focused on the measurement of the added worker effect, i.e. married women's tendency to increase their labor supply when their spouse is hit by a negative income shock (Lundberg, 1985). Recent contributions include Halla et al. (2020) who investigate married woman's labor supply decisions in response to their husbands' job displacement in Austria; they find evidence of a small added worker effect, which amounts to a 1 percentage point increase in the employment rate. While our estimates are not significantly different from 0, they are compatible with such tiny increases, even in a context in which the baseline female employment rate is much higher.¹ With respect to this paper, we extend the methodology to investigate men's response to their female partners' job displacement, and find similar results. By contrast, other recent studies, more structural on their approach, tend to find sizable spousal labor supply responses to individual wage shocks (Blundell et al., 2016). More broadly, our results suggest that couples react very little to shocks that affect the within-family comparative advantage (Becker, 1981): otherwise, job displacement that decreases the returns on time and effort spent by one family member on the labor market, keeping her home-production productivity constant, should induce the reallocation of time and effort of her spouse towards the labor market.

A key element to explain these differences in family insurance mechanisms may be related to the collective insurance offered both by unemployment insurance and

¹The employment rate for male displaced workers' spouses is about 46% before job displacement in Austria, against more than 80% in the French setting.

the progressive taxes and transfers system, if the latter crowd out the former (Cullen and Gruber, 2000; Autor et al., 2019b). Our third contribution thus relates to the role played by these safety nets. Indeed, we show that in the short run, unemployment benefits provide quite effective insurance against the income shock that results from job displacement: its magnitude would be two to three times larger absent unemployment insurance. Yet unemployment insurance offers much less coverage against the long-run component of the shock, which is not driven by employment. Workers seem partially insured at best against the long-run component of the shock, through a bundle that involves not only self-employment and capital markets, but also the progressivity of the taxes and transfers system, which is consistent with recent results based on Canadian data (Stepner, 2019).

The remainder of the paper is organized as follows. Next section presents the data and describes the constitution of our sample of displaced and control workers. Section 4.3 details our empirical framework. Section 4.4 presents the results and Section 4.5 concludes.

4.2 Data

Our analysis relies on two French administrative datasets. Firstly, we rely on the *Déclaration Annuelles de Données Sociales* (DADS), which consist in comprehensive payroll tax registers that cover the universe of French salaried employees, to identify sharp decreases in plant size that we interpret as mass layoffs.² Secondly, we use these mass layoffs to identify the consequences of exogenous job separations in the *Échantillon Démographique Permanent* (permanent demographic sample, EDP), a combination of administrative registers that links payroll tax data with tax returns and birth registers, thanks to a common identifier based on a Social security number.³

²Mass layoffs are strictly regulated under French law. In Appendix 4.A, we detail the most salient features of this regulation. In practice and due to data limitations, our empirical approach may depart from the precise, legal definition of a mass layoff.

³In practice, tax returns are linked with other datasets based on variables with high identifying power (name, surname, gender, date of birth, birth location).

4.2.1 Payroll tax registers

Our labor market data are drawn from the *Déclarations Annuelles de Données Sociales* (DADS). By law,⁴ French employers have to fill in a DADS form for every employee subject to payroll taxes. The form contains detailed information about days paid, hours paid, occupation, industry, gross and net wages, other job characteristics (beginning, duration and end of a period of employment and part-time employment), employer characteristics (size and location) and individual characteristics (age, gender and municipality of residence). These forms are gathered in a linked employer-employee database: plants are identified by the Siret, a 14-digits plant identifier, and individuals are identified by a anonymized identifier based on the NIR, a Social security number.⁵

4.2.2 Permanent Demographic Sample

The *Échantillon Démographique Permanent* (permanent demographic sample, EDP) gathers administrative information from birth registers, tax returns, firm records and from the file of welfare benefits recipients, for a representative sample of the French population including all individuals born one of the first four days of April, July, October or between January 2nd and 5th. The EDP follows the professional, financial, residential and family situation of these "EDP individuals" each year between 2010 and 2016. On top of this core sample, for each year, the EDP features information on each person listed in the same dwelling.⁶

Key to our empirical approach is the fact that the EDP data include an extract of the DADS data. This allows us to link these data to the comprehensive payroll tax data, and specifically to determine treatment and control groups based on plant-level in- and outflows that can be observed in the comprehensive DADS registers.⁷ We then consider the consequences of these exogenous employment shocks in terms

⁴The absence of DADS as well as incorrect or missing answers are punished with fines.

⁵In practice, this anonymized identifier can only be used to track individuals over two consecutive year of the data.

⁶Precisely, information is available for each member of a tax household listed in this dwelling, including in particular young adults living apart from their parents but still registered on their income tax return and excluding people registered in another tax household (for instance a student from another family who is renting a room there). Individuals living in collective housing and homeless people are excluded from the EDP.

⁷In practice, the plant-level identifier is not available in the extract of the DADS that is included in the EDP data. Instead, we link the two datasets based on individual and job-level variables with high identifying power. The accuracy of this statistical matching approach is over 98%.

of family structure and income, based on income tax returns and birth registers that form part of the EDP data.

Tax returns

Household structure The EDP data provide information based on income tax returns. This information is available at (i) the individual level for some income sources, especially labor earnings or unemployment benefits, and (ii) at the household level for all other income sources, including capital income or social benefits. In other words, the structure of the data assumes perfect pooling of capital income within the household.

This household level is not defined by fiscal law: it is a statistical concept which the definition is managed by Insee during the recollection of tax returns data. The definition of this household level is based on the pairing of income tax returns and housing tax returns. In the context of this particular paper, this approach is highly desirable because it allows us to treat cohabiting and married couples the exact same way, even though the former fill in separate income tax returns while the latter fill in joint income tax returns as married couples are taxed jointly. This is particularly important in a context in which marriage is increasingly delayed or abandoned in favor of civil unions and cohabiting relationships: in 2014, slightly more than one cohabiting heterosexual couples in three had married during the first 5 years of co-residency ([Costemalle, 2015](#)).

On top of delineating households, Insee provides variables that describe the type of household at stake. These variables are based on (i) the family information that appears in the tax returns, and (ii) the age and gender of the members of the household. Indeed, especially in the case of unmarried couples, tax returns alone are not sufficient to assess the type of relationship that may exist between cohabiting adults. Specifically, two different-sex cohabiting adults are considered to be in a heterosexual cohabiting relationship if either (i) they appear as married, have signed a civil union contract or declared themselves to be a couple in their housing tax return; or (ii) they live together and have an age difference of 15 years or less. Based on age differences, Insee then assesses whether other inhabitants of the household are likely to be children of a couple or of one of the older inhabitants; this is relatively easy in the case of minor children who appear as such on income tax returns, since they allow their parents to benefit from a tax rebate, but can

get complicated for older children. For this reason, and also due to a left-censoring issue, we base ourselves on birth registers when investigating fertility decisions.

In the end, we rely on this imputed information in two ways. Firstly, all our estimates are conditional on household type as measured before job displacement. This household type is a categorical variable with 5 categories: single individuals, childless couples, single-parent families, couples with children and lastly a remaining category which gathers all other situations. Secondly, we create a dummy variable that equals 1 if a displaced (control) worker is either the main registrant of the household or the spouse of the main registrant of the household, and 0 otherwise, which allows us to observe couple formation and dissolution in response to job displacement. This approach is unsuited for same-sex couples, which is why, on top of sample size considerations, we restrict ourselves to the study of heterosexual relationships.

Income variables The income variables extracted from the income tax returns are available at both the individual and the household level. Our analysis relies on both levels. Firstly, we consider the individual income data related to displaced workers or their spouses. This individual income sums wage income,⁸ unemployment benefits, self-employed earnings, pensions and alimonies. Throughout the paper, we aggregate the last three components together. Our sample excludes workers that would be eligible for a retirement pension after job loss. Additionally, the payment of alimonies is triggered by couple dissolution, that we do not find to increase in response to job displacement. As a result, the effect of job displacement on this component corresponds to adjustments through self-employment. These income variables mix up all margins of labor supply, i.e. whether to work or not, but also how many hours, for which kind of employer, the level of effort etc. Most of these margins cannot be disentangled with the data. However, we can get a sense of the extensive margin of employment by considering three dummy variables: (i) having positive wage income (salaried employment); (ii) having positive wage income or positive self-employed earnings (employment); (iii) having positive wage income or positive self-employed earnings or positive unemployment benefits (participation).

⁸The wage income concept in terms of the income tax differs from the wage income concept in terms of the payroll tax. For instance, sick leave and maternity leave benefits are part of the wage income in terms of the income tax, but are excluded in terms of the payroll tax. These differences may to some extent explain differences between our wage income estimates and those of other recent papers that rely on the DADS data, such as [Brandily et al. \(2020\)](#).

Secondly, we rely on the household-level income data. These data allow us to consider additional margins of adjustment that cannot be observed in the individual data. Specifically, our focus is on household disposable income, which equals the sum of all individual incomes described *supra* plus capital income (excluding capital gains) and family and welfare benefits, less housing, property and income taxes. This allows us to consider not only the role played by unemployment benefits in providing insurance against these large income and employment shocks, but also whether households are able to effectively insure themselves thanks to the capital market – which would show up in the capital income component – and the insurance value of the progressive taxes and transfers system.

The income data is sometimes missing: empirically, for each year of the data about 8% of displaced and control workers cannot be observed in the income tax returns. In practice, our estimates are conditional on being observed in the income tax returns data, but not necessarily continuously so. In Appendix 4.E, we assess the robustness of our results with respect to this choice by (i) displaying estimates of the effect of job displacement on the probability of being present in the income tax returns data; and (ii) by replicating our analysis on a sample of workers who are continuously observed in the data. The sample inclusion rate decreases very little due to job displacement. However, our results turn out to be immune to this issue.

Our data include negative incomes, as self-employed workers are allowed to report losses in their income tax returns. These negative incomes are quite uncommon; however, we choose to drop observations related to households for which negative incomes are observed even once. Our approach investigates income measured in levels, as opposed for instance to the logarithm of income. This allows us to include individuals or households whose earnings equal 0. However, this could also make our estimates less robust with respects to outliers in the right tail. For this reason, we winsorize all earnings and income variables at the 99th percentile level.

Birth registers

Our analysis also relies on birth registers which form part of the EDP data. Births are registered by an individual who was present at the time of birth, usually the father, but in some cases a doctor or a midwife. These registers allow us to follow the fertility decisions of EDP individuals, since they cover children of EDP individuals. Specifically, children born to EDP individuals may be tracked from 1968 when their

EDP parents are born on October, and 2004 for the others. This creates a left-censoring pattern in the data. For this reason, we focus on the number of children born each year around the employment shock, as opposed to the overall number of children born to EDP parents,⁹ which would incorporate unobserved past fertility decisions.

Birth registers in the EDP incorporate three distinct types of events: usual childbirths, still births and adoptions. For the sake of this particular paper, we choose to consider the first two types of events, that will capture (potential) parents' decision to have children, regardless of the outcome. The criteria according to which still births are registered varied over time. Since 2008, still births are registered on the basis of a medical certificate of childbirth (delivery). However, the still birth rate is sufficiently low for this particular choice not to affect our results.¹⁰

4.2.3 Job displacement identification and sample construction

We identify exogenous shocks on the employment relationship of an individual based on in- and outflows measured at the plant-level in the DADS data. Let T denote a year between 2011 and 2017. We begin by selecting private-sector plants (i) with more than 10 workers on January, 1st $T - 1$; (ii) of which at least 25% of workers present on January, 1st $T - 1$ left the plant during year $T - 1$; and (iii) of which the number of workers decreased by at least 25% between January, 1st $T - 1$ and December, 31st T . Condition (iii) ensures that we are focusing on large decreases in the workforce of a particular plant, as opposed to plants in which the turn-over is extensive. To ensure that these events do indeed correspond to mass layoffs, as opposed to mere changes in plant-identifier or a spin-off, we then track workers who left the plant during year $T - 1$ and recover their employers during year T . We impose that no more than 25% of the workers who left the plant work in the same plant during year T ; this kind of restriction can be found for instance in [Gathmann et al. \(2020\)](#). Lastly, when a plant matches these criteria for multiple years, we only keep the first event.

Armed with this collection of plant identifiers affected by mass layoffs, we turn to the individual-level DADS data that form part of the EDP sample. We consider

⁹The former being the derivative of the latter.

¹⁰The still birth rate is about 1% since this law change.

an individual to be affected by a mass layoff if (i) she works in an affected plant and leaves it at some point between one year before and one year after the plant-level event, (ii) her tenure within that plant by the time she leaves is at least 3 years and (iii) she is between 25 and 50 years old by the time she leaves the plant. Condition (ii) ensures that we do not consider temporary workers for which job separation cannot be regarded as a shock; it further implies that the separations upon which we focus do not correspond to young workers' gradual entrance into the workforce, that usually involves multiple job transitions (Topel and Ward, 1992). Condition (iii) ensures that we are not dealing with individuals for which retirement is an option, so as to simplify the interpretation of our results.

Our empirical analysis requires that we compare these *displaced* workers with a set of *control* workers to quantify the impact of these employment shocks. To determine the set of control workers, we consider workers (i) who work in the same detailed (5-digits) industry and meet the same tenure requirement in their own plant at the time of the (counterfactual) job separation and (ii) have never worked in one of the affected plants. A similar individual may appear as a relevant control for multiple events; in such a case, we sample the timing of her counterfactual shock among all those for which she appears as a possible control with a uniform probability. To ensure that we have sufficiently large sample sizes and to support the assumptions upon which our empirical analysis rests, we drop detailed industry \times timing of the shock cells (i) that contain 5 controls individuals or less, or (ii) in which more than half of the individuals are treated. Condition (i) sustains the credibility of the common support assumption that motivates our reweighting approach. Condition (ii) rules out industry-level shocks: affected plants cannot gather the vast majority of the workforce in their industry. This makes it less likely that other plants of the same industry are affected by their hiring and laying-off decisions. It also makes it more plausible that control workers' job opportunities are left unaffected by the shock, which is crucial for the implicit Stable Unit Value Treatment Assumption upon which our empirical framework rests.

Due to data limitations and the fact that we allow for some anticipation in the effect of job displacement, we only focus on individuals that are affected by a (counterfactual) shock between 2012 and 2016. This leaves us with over 315,000 individuals who we observe between 2010 and 2016 in the income tax returns data, and between 2004 and 2016 in the birth registers.

4.2.4 Summary statistics

Figure 4.1 displays the distribution of the sizes of the affected plants, compared to that from which control workers originate. These sizes are measured as the number of workers on January, 1st, one year before the (counterfactual) plant-level event, and the distribution are weighted by the number of workers observed in our sample. Many workers in our sample originate from relatively small plants. However, quite reassuringly, this is also the case in our control group. This suggests that our delineation of the displaced group based on observed plant-level outflows does not lead us to focus on spurious shocks that mostly stem from small sample sizes.

Figure 4.2 displays the distribution of the intensity of the plant-level shock in our displaced group. This intensity is measured by one minus the size of the plant on December, 31st one year after the shock, relative to this size on January, 1st almost two years before. Strikingly, over half of our displaced worker group were affected plant closure events.

Table 4.1 compares displaced and control workers, separately by gender, in terms of their ages, their occupation before the (counterfactual) shock and their family structure before the said shock. Regardless of gender and job displacement exposure, the average age in our sample is about 37. Even though there are some small differences in terms of occupation or family structure, the differences between displaced and control workers with respect to these variables, as measured two years before the (counterfactual) shock remain limited. Most workers live with a partner: two years before they were displaced, over 7 out of 10 displaced workers lived with a partner, almost always with a different-sex partner. The proportion is very close when it comes to our control group of non-displaced workers. Similarly, over 60% of displaced workers lived with children, and the same goes on for control workers.

Table 4.2 compares workers in terms of their labor earnings as observed in the payroll tax data, as well as their hours and days worked, to get a sense of their labor force attachment. All these variables are collected two years before the (counterfactual) job loss, so as to represent the baseline levels from which workers will evolve due to job separation. Control and displaced workers appear quite comparable along these dimensions. Prior to job displacement, workers earned about €20,000 for women and €25,000 for men. Their days worked (which include paid vacations and week-ends) amount to 340 days a year, which suggest a substantial labor force attachment, since in the DADS data an entire year of employment corresponds to

360 day. Lastly, when (i) dividing by the average days worked and (ii) comparing it to the baseline level for a full-time job in the DADS data (1,820 hours a year), the average hours worked imply that women work about 90% of a full-time equivalent in average, whereas men are almost always on a full-time basis.

We then compare displaced and control workers two years before the shock in terms of their income, as measured in the income tax returns data. We begin by workers' own individual earnings, as displayed in Table 4.3. Differences between displaced and control workers remain limited in that matter. Before displacement, displaced workers earn about €22,000 for women and €27,000 for men. These earnings mostly result from wage income, as opposed to other sources of income. Wage income observed in the income tax returns is slightly higher than wage income as observed in the payroll tax data (see Table 4.2); this is mostly due to the fact that the latter is more restrictive, whereas the former includes sources of earnings that are not considered as wages in terms of the payroll tax, e.g. sick leave benefits or maternal leave benefits. We check that consistent with our framework, almost every worker is observed in employment in the income tax returns data: the employment rate is very close to 100%, which speaks in favor of the consistency of income tax returns data with payroll tax data filled by employers.

We replicate this exercise, this time considering displaced and control workers' spouses' earnings before the (counterfactual) shock, thus restricting to the 70% of workers who live with a partner before the said job separation. Table 4.4 displays our results. Here again, differences between displaced and control workers are not very large. Consistent with typical gender inequality, displaced women's male partners earn much more than displaced men's female partners. Here again, earnings are driven by wage income as opposed to other sources of earnings. The participation and employment rates for women's male partners are 95%, whereas they are between 80% and 85% for men's female partners. These rates match quite closely the participation rate for workers aged 25 to 49 in France during the relevant time period. Due to the gender gap in participation and employment, there is less room for positive labor supply decisions at the extensive margin for men's female partners than for women's male partners, which induces most papers devoted to the added worker effect to focus on displaced male workers' female spouses.

Lastly, we consider income at the household level, as measured two years before the (counterfactual) job separation. Table 4.5 displays our comparison. Here again, differences across treatment groups are quite small. The average household

disposable income ranges between €45,000 and €47,000 before job separation. This disposable income mostly results from wage income: other sources of earned and non-earned income matter much less. The average household pays between €7,000 and €8,000 a year in taxes (which sums income, housing and property taxes). We then consider the position of households in the equivalent income distribution, that is the distribution of disposable income divided by the number of consumption units in the household. Our sample overrepresents the middle of the distribution, and underrepresents the bottom of the distribution.

4.3 Empirical analysis

Our empirical analysis leverages differences in the evolution of family structure and labor outcomes of treated and control over time, after conditioning on relevant observables. To this end, we rely on a reweighted difference-in-difference approach developed by [Abadie \(2005\)](#). This section details this framework, and presents a way to quantify the implications of a violation of the parallel trends assumption upon which it relies when it comes to spousal labor supply.

4.3.1 Identifying the consequences of job displacement

Let $Y_{i,t}$ denote a variable that describes individual i 's family structure or income at time t , which is measured relative to the year during which her (counterfactual) employment shock occurs. Let D_i be a dummy variable that equals 1 if individual i belongs to the set of displaced workers, and 0 otherwise, in which case she belongs to the control group. Let C_i denote her cohort, which corresponds to the year during which her (counterfactual) job displacement takes place. Lastly, let X_i be a vector of time-constant observable variables, which includes: the industry in which worker i worked before the (counterfactual) employment shock, her 1-digit occupational group, her labor market experience and her tenure within the firm by the time she is displaced, her year of birth, and her family structure and the income quintile to which she belongs two years before the shock. In all what follows, all expectations and probabilities are taken conditional on gender; we keep this conditioning implicit throughout this subsection.

We define $Y_{i,t}(d)$ to be the potential outcome of individual i at time t , depending on whether she has ($d = 1$) or she has not ($d = 0$) been displaced. We are interested

in the causal effect of job displacement on Y , i.e. (functionals of) the distribution $Y_{i,t}(1) - Y_{i,t}(0)$. Specifically, we consider the cohort-specific average treatment effect on the treated:

$$CATT(c, t) = \mathbb{E}[Y_{i,t}(1) - Y_{i,t}(0) \mid C_i = c, D_i = 1] \quad (4.1)$$

Our identification strategy is based on three assumptions:

Assumption 7 (Common support). *For all c , for all x :*

$$0 < \mathbb{P}(D_i = 1 \mid C_i = c, X_i = x) < 1 \quad (4.2)$$

Assumption 8 (Parallel trends in baseline outcome). *For all c , for all x , for all t, t' :*

$$\begin{aligned} & \mathbb{E}[Y_{i,t'}(0) - Y_{i,t}(0) \mid C_i = c, X_i = x, D_i = 1] \\ &= \mathbb{E}[Y_{i,t'}(0) - Y_{i,t}(0) \mid C_i = c, X_i = x, D_i = 0] \end{aligned} \quad (4.3)$$

Assumption 9 (Limited anticipation). *For all c , for all x , for all t , if $t < -1$, then:*

$$\mathbb{E}[Y_{i,t}(1) - Y_{i,t}(0) \mid C_i = c, X_i = x, D_i = 1] = 0 \quad (4.4)$$

Assumption 7 states that every worker affected by an employment shock has at least one non-affected surrogate worker with the exact same observable characteristics. Assumption 8 states that absent the employment shock, affected workers' average outcomes would have evolved the same as their non-affected surrogates' average outcomes. Lastly, Assumption 9 states that the average effect of job displacement equals 0 up until two years before the shock. The effect is not constrained to 0 one year before the actual job separation, so as to allow workers and household to anticipate these shocks that are not perfectly unforeseeable.

Under these assumptions, provided that individuals can be observed sufficiently long after job displacement, cohort-specific ATTs can be identified from the data:

Proposition 3 (Difference-in-difference estimand). *Let $\{\underline{T}, \underline{T} + 1, \dots, \bar{T} - 1, \bar{T}\}$ denote the set of years that can be observed in the data. For c and t such that $\underline{T} + 1 < c < \bar{T} + 1$ and $\underline{T} - 1 < c + t < \bar{T} + 1$, $CATT(c, t)$ can be identified from the*

data and:

$$\begin{aligned}
CATT(c, t) &= \mathbb{E}[Y_{i,t} \mid C_i = c, D_i = 1] \\
&\quad - \mathbb{E}[Y_{i,c-2} \mid C_i = c, D_i = 1] \\
&\quad - \mathbb{E}[\pi(c, X_i)Y_{i,t} \mid C_i = c, D_1 = 0] \\
&\quad + \mathbb{E}[\pi(c, X_i)Y_{i,c-2} \mid C_i = c, D_1 = 0]
\end{aligned} \tag{4.5}$$

where $\pi(c, x) = \frac{\mathbb{P}(D_i=1 \mid C_i=c, X_i=x)}{1 - \mathbb{P}(D_i=1 \mid C_i=c, X_i=x)} \frac{1 - \mathbb{P}(D_i=1 \mid C_i=c)}{\mathbb{P}(D_i=1 \mid C_i=c)}$

In other words, it is possible to reweight the data so that the distribution of covariates in the control group matches that of the treated group. Because the parallel trends holds conditional on these covariates, as long as both the pre-treatment period and the post-treatment period are observed, it is possible to implement usual difference-in-difference techniques on the reweighted data. Our summary statistics show that displaced and control workers are very similar before the shock: our reweighted difference-in-difference approach makes them even more comparable so as to achieve a credible identification of the causal effect of job displacement.

Proposition 3 allows us to identify cohort-specific ATTs. Our aim is to aggregate them across cohorts to gain insights into the dynamic effects of job displacement. To do so, for each t , let $\mathcal{C}(t)$ denote the subset of cohorts for which $CATT(c, t)$ can be identified from the data.¹¹ We consider the aggregate average treatment effect on the treated:

$$\Delta(t) = \mathbb{E}[Y_{i,t}(1) - Y_{i,t}(0) \mid D_i = 1, C_i \in \mathcal{C}(t)] \tag{4.6}$$

This quantity represents the average difference between the realized outcomes and the counterfactual situation in which workers would not have been displaced, t years after the job separation took place, for workers affected by shocks of which the consequences can be inferred from the data. Because the data only cover a restricted time-period, this set of workers varies depending on t . By the law of iterated expectations:

$$\Delta(t) = \sum_{c \in \mathcal{C}(t)} \mathbb{P}(C_i = c \mid c \in \mathcal{C}(t)) CATT(c, t) \tag{4.7}$$

¹¹Specifically, $\mathcal{C}(t) = \{\underline{T} + 2, \underline{T} + 3, \dots, \bar{T} - 1, \bar{T}\} \cap \{\underline{T} - t, \underline{T} - t + 1, \dots, \bar{T} - t - 1, \bar{T} - t\}$

By Proposition 3, as long as $\mathcal{C}(t)$ is non-empty, it is therefore possible to express $\Delta(t)$ in terms of quantities that are all identified from the data. The non-emptiness condition is equivalent to $t \in \{\underline{T} - \bar{T}, \underline{T} - \bar{T} + 1, \dots, \bar{T} - \underline{T} - 3, \bar{T} - \underline{T} - 2\}$. With $\underline{T} = 2010$ and $\bar{T} = 2016$ we are therefore able to cover dynamic effects that span from six years before job displacement to four years after.

Combined with Proposition 3, Equation 4.7 suggests a very simple choice of plug-in estimator, in which we substitute expectations and probabilities with their empirical analogues. In practice, we base our reweighting on conditional probabilities drawn from a probit model estimated by maximum likelihood. Under standard regularity and integrability assumption, this estimator is asymptotically normal. To conduct inference, we rely on a bootstrap approach, clustered at the level of the plant at which workers were working immediately before the (counterfactual) employment shock hit them. This choice is based on the premise that this is the level at which shocks are assigned (Abadie et al., 2017), given that we rely on in- and outflows measured at the plant-level to identify these shocks.

4.3.2 Quantifying the discouraged worker bias

A pitfall of our approach when it comes to spousal labor supply decisions is that couples may be affected by positively correlated shocks that can lead to violations of our conditional parallel trends assumption. Specifically, the EDP sampling plan is defined *at the individual level*. Consistently, this is also the level at which we determine our treated and control groups. If individuals are affected by both shocks to their own employment status, and shocks to their spouse's employment status, and if these shocks are positively correlated, the problem is thus that, for individuals of which we do not observe the employment status in the DADS data, we are likely to interpret as the consequences of a shock to the employment status of their spouse what actually results from shocks to their own employment status. This positive correlation between employment shocks of spouses has been referred to as the *discouraged worker effect*, as opposed to the added worker effect. It is especially pronounced in settings in which spouses actually evolve in the same labor market, which is plausible for instance given that 10% to 20% of couples meet at work (Bozon and Rault, 2012; Rosenfeld and Thomas, 2012), and would result in downward bias in our estimation of the effect of the labor supply effect of a shock on one's spouse employment status. In other words, we would conflate changes in spousal

labor supply with changes in the labor demand for the work of displaced workers' spouses.

To quantify the bias that results from this violation, in Appendix 4.B, we develop a simple model of within-couple correlated employment shocks that allows us to approximate the consequences of this departure from our identifying assumptions. Under additional assumptions that (i) the effect a shock on a worker's employment status on her labor outcomes dominates the effect of a shock on her spouse's employment status; (ii) the within couple correlation of employment shocks remains limited; (iii) the direct effects of a shock on one's own employment status do not depend on whether one's spouse is affected by a shock or not; (iv) the only channel that generates the within couple correlation in employment shock is that spouses work in the same firm (detailed industry); and (v) the probability of receiving an employment shock does not depend on whether or not one works in the same firm (detailed industry) as her partner, we show that the bias in women's (men's) reaction to a shock on their male (female) partner's employment status is well approximated by the direct effect of a shock on women's employment status on women's outcomes, multiplied by the share of couples that work in the same firm (detailed industry). The first quantity is a direct byproduct of our identification strategy. A reasonable estimate of the second term can be obtained from the French Labor Force Survey.

4.4 Results

4.4.1 Couple formation and dissolution

Figure 4.3 displays our estimates of changes in displaced workers' probability of living in a heterosexual cohabiting relationship, relative to the (reweighted) control group, over time relative to job separation, and separately by gender, for workers who lived with a spouse two years before (counterfactual) job displacement. Table 4.1 shows that this group gathers about three quarters of our sample. Under Assumptions 7 to 9, these quantities can be interpreted as the average dynamic causal effect of job displacement on couple dissolution.

Firstly, we find almost no difference in the evolution of the probability of being in a heterosexual cohabiting relationship before job displacement: the difference in trends before job separation is close to 0. In other words, displaced and control workers evolve in parallel in that matter. This fact strengthens the credibility of

our empirical framework, as our model predicts that the causal effect of job displacement is 0 up until two years before the separation actually takes place. While this is not sufficient to assess the plausibility of our identifying assumptions *after* job displacement takes place, it does suggest that a causal interpretation of these quantities is possible.

Secondly, our estimates regarding the post-job separation period do not suggest that job separation disrupts couples. Indeed, our estimates remain precise and centered on the 0 line when it comes to displaced women, which suggests that the probability of being in a heterosexual relationship does not decrease due to job separation. When it comes to displaced men, our estimates are compatible with a very small decrease over time due to job displacement. The magnitude of this decrease, about 2.5 percentage points four years after an exogenous job separation is very small, and we can reject decreases larger than 5 percentage points over the same time period. Furthermore, the width of our confidence intervals is still compatible with null effects.

We then replicate this exercise, this time restricting ourselves to workers *who did not live with a spouse two years before job displacement*. This group gathers about a quarter of our sample. Under Assumptions 7 to 9, the estimated quantities would correspond to the dynamic effect of job displacement on couple formation, i.e. whether job displacement increases or decreases the probability of finding a mate for those who did not have one. Figure 4.4 displays our results.

Here again, we find very limited differences between displaced and control workers in terms of the evolution of their outcomes before the shock. In other words, displaced and control workers evolve in parallel over time before job displacement occurs, which helps sustain the validity of our identifying assumptions. After job displacement, the probability of living in a heterosexual cohabiting relationship seems to increase for women, but not for men. However, the magnitude of our estimate for women – 6 percentage points 4 years after job displacement – remains small, and our confidence intervals are still compatible with null effects. Lastly, we do not find formerly single men to change their probability of living in a cohabiting heterosexual relationship due to job displacement.

4.4.2 Fertility decisions

Figure 4.5 plots changes in the number of children born to displaced workers over time relative to job loss and relative to the (reweighted) control group, separately by gender. Here again, we find that displaced and control workers evolve in parallel over time *before* job separation, which supports the credibility of the identifying assumptions under which these estimates have a causal interpretation. We do not find any change past job separation, which suggests that job displacement does not lead to changes in fertility decisions.

4.4.3 Labor supply, taxes and transfers

We now turn to the consequences of job displacement on earnings, labor supply and transfers, first focusing on the displaced worker, then on her spouse's labor supply, and lastly on the household as whole, which will allow us to assess the role played by family insurance and the taxes and transfers system in insuring workers against exogenous shocks to the employment relationship.

Own earnings and labor supply

Figure 4.6 displays our estimates of changes in displaced workers' individual earnings and labor supply, as observed in the tax returns data, over time and relative to the (reweighted) control group. Displaced and control workers' earnings and labor supply evolve roughly in parallel before the shock.

Under our identifying assumptions, our estimates imply that, consistent with the intuition, job displacement results in a steep drop in workers' wage income. In the short run, the magnitude of this drop is about €5,000 for women and €6,300 for men. This short run transitory wage income shock is likely tied to the drop in employment: the yearly employment rate, as measured in the tax returns data by delineating the group of workers with positive labor earnings, drop by 7 (for women) to 8.5 (for men) percentage points.

However, in the longer run, the employment effect of job displacement fades out, which suggests that most displaced workers manage to find new jobs on the labor market at the same rate as their control counterpart. By contrast, while the magnitude negative wage income shock decreases, it remains notable, corresponding to 10% to 15% of the baseline pre-shock wage income (see Table 4.3). This suggests

that this long run earnings drop is not driven by the displaced workers' inability to find new jobs, but rather by other channels such as the dissolution of valuable employer-employee matches (Lachowska et al., 2020). In Appendix 4.C, we replicate this analysis, relying on payroll tax data issued from the DADS instead of tax returns, to assess the robustness of these results. Our estimates are similar in the long run, and differences in the short run are likely the result of differences between the conceptual measures of outcomes in payroll tax data and income tax returns.

While the short run drop in wage income is very steep, job displacement does not result in similarly massive losses in individual earnings. The reason for this is that, consistent with the rationale, unemployment benefits limit the loss in earnings. Specifically, the insurance role of unemployment benefits allow workers to move from a €5,000 to a €1,800 short-run loss when it comes to women, and from a €6,300 to a €3,100 short-run loss for men. In other words, unemployment insurance divides the short-run earnings loss by a factor 2 to 2.8. Other sources of individual earnings, which consist of self-employed earnings, pensions and alimony matter much less in the short run. This is not surprising given that (i) workers in our sample are too young to be eligible for retirement pensions, (ii) moving from a salaried job to self-employed employment takes some time and (iii) alimonies are tied to couple dissolution, which we did not found to be triggered by job displacement at least in the short run.

This is no longer true in the longer run, which is not surprising given that the negative employment effects fade out over time, which will mechanically decrease the relevance of unemployment insurance. Four years after job displacement, we can no longer reject the null hypothesis that the overall individual earnings loss is the same as the wage income loss. Our baseline estimates would imply that other sources of income decrease the wage income loss by about 25% for women and 33% for men. Even so, while job displacement does not result in long-run differences in the employment rate, it does cause long-run earnings loss, of which the magnitude is about the same as the short-run earnings loss. This would implies that, *in terms of their individual earnings*, while workers are well insured against transitory employment shock thanks to unemployment insurance, there are much less insured against permanent wage shocks.

Spousal earnings and labor supply

Figure 4.7 displays our raw estimates of changes in displaced workers' spouses' earnings and labor supply, over time and relative to control workers' spouses. This estimation can only be performed when the displaced worker is observed to have a different-sex spouse.¹² Given that we do not find job displacement to change the probability of living in a cohabiting heterosexual relationship (Figures 4.3 and 4.4), this should not generate spurious patterns due to sample selection.¹³ We find very little differences between displaced and control workers' spouses' trends in earnings and labor supply. Specifically, we cannot reject the null hypothesis that displaced and control workers' spouses' earnings and labor supply evolve in parallel over time both before and after job displacement.

Discouraged worker effect As noted in Section 4.3, these results might be biased due to the within-couple correlation in employment shocks. To circumvent this issue, we showed that under simple assumptions regarding the structure of these within-couple correlations, and reasonable approximations, it is actually possible to implement a correction for the resulting downward bias in spousal labor supply. Specifically, we showed that this bias is well approximated by the product of the direct effect of one's own displacement on one's own labor supply and earnings by the share of couples who work for the same employer.

Based on the French Labor Force Survey (*Enquête Emploi en Continu*), we estimate that between 2010 and 2012, about 5% of men (women) who lived in a cohabiting relationship with a women (men) aged between 25 and 50 who held a job as a salaried employee worked in the same plant as her (him). The corresponding shares when it comes to working in the same industry are 8% and 6%.

We combine these shares with our estimates of the direct effect of job displacement to correct the bias that results from the within-couple correlation of employment shocks in our investigation of the added-worker effect. Figures 4.8 and 4.9 display our results. Because the share of couples for which employment shocks are

¹²Displaced workers' individual loss are broadly the same regardless of whether they are with spouse or not. In Appendix 4.D we replicate Figure 4.6 while conditioning on being observed with a spouse. Our estimates appear extremely similar to our baseline results.

¹³To assess the robustness of our conclusions with respect to this concern, in Appendix 4.E.3, we replicate our analysis on a balanced panel of workers who are continuously observed with a spouse. Our results are very similar to our baseline estimates, which confirms that this sample selection is not an issue here.

correlated is actually quite low, these results are very similar to those obtained without this correction. In other words, the added worker effect is not active: displaced workers' spouses do not increase their labor supply to compensate for the income loss that results from job displacement.

Total household income

Figure 4.10 displays changes in all the layers that range from wage income to disposable income for displaced workers' households, over time and relative to control workers' households. All components of disposable income evolve roughly in parallel between the displaced workers' group and the control group, which makes our identifying assumption reasonable. Crucially, income tax returns allow us to go through all the layers of both the taxes and transfers system, and earnings sources that combine into household's disposable income. As a result, we are able to get a sense of which sources of income provide household some kind of insurance against the income risk that job displacement generates.

Consistent with the rationale and our previous results, households experience sharp drops in wage income due to job displacement in the short run. The magnitude of these wage income losses coincides with the sum of the own wage income reaction and that of the spouse (which is close to 0). Once again as the individual-level data suggests, unemployment insurance provides effective insurance in the short run: unemployment benefits decrease the short-run wage income loss by more than 60%. Other components seem to matter much less in the short run, which may stem from a variety of frictions.

In the longer run, unemployment benefits seem to matter much less. This is consistent with our previous results, and the fact that the employment effect of job displacement is only short-lived, which mechanically limits the relevance of unemployment insurance. While our estimates get noisy over time because we rely on a more restricted subsample, they still suggest that (i) job displacement results in a persistent shock on households' disposable income, which remains approximately constant over time, by contrast with the wage income component; and (ii) that households are partially insured at best against the long run wage income risk through a bundle that involves not only unemployment insurance, but also in a somewhat similar extent self-employment, capital income, and lastly the taxes and transfers system. This partial insurance does not seem to be active when it comes to

displaced women. Overall, when compared with the pre-job loss baseline, this persistent disposable income loss amounts to a 3 to 6% decrease in household income, regardless of whether we consider the short-run or the long-run effects. Even though the data show that households are only partially insured against such shocks, the contrast is striking with the short-run individual wage income loss, which amounts to a 25% drop.

4.4.4 Selection into job displacement

Job displacement is not randomly assigned to workers: it results from decisions that may be endogenous with respect to the outcomes we investigate. Firstly, with the exception of plant closures, displaced workers are a subset of the workforce of the firm, which is likely to be chosen at least partly based on their contribution to the firm's product. For instance, [Seim \(2019\)](#) shows that within Swedish firms, workers who display high cognitive and noncognitive skills are less likely to be displaced than their counterparts with poorer skills. Secondly, when it comes to mass layoffs and even more plant closures, the assignment of job displacement is actually dependent on workers' past labor supply decisions: those with the best outside options may leave the firm *before* the decision to dismiss part of the workforce is made, solely based on their knowledge of the economic shock that triggers the layoff decision. The first channel is particularly salient in the French case: by law, employers are bound to either resort to criteria defined by industry- or firm-level collective agreements, or to define their own criteria, that always have to include: (i) workers' families (especially when it comes to single parents); (ii) workers' tenure within the firm; (iii) workers' employment prospects (especially when it comes to workers with disabilities or older workers); (iv) and lastly workers' productivity (see [Appendix 4.A](#)).

To assess the robustness of our results with respect to these issues, we consider several quantitative exercises. Firstly, we replicate our framework while restricting the treatment group to the subset of workers who were affected by plant closures. These workers constitute over half of our baseline treatment group (see [Figure 4.2](#)). Because all workers of the firm are displaced, here the selection into displacement should be less active than it is otherwise. [Appendix 4.F.1](#) displays our results, which are all consistent with our baseline estimates.

Secondly, within affected plants, we compare displaced and non-displaced workers. To do so, we restrict ourselves to plants that are affected by a mass layoff, but

not by plant closure: otherwise the non-displaced group is empty. We replicate all our summary statistics tables to get a sense of how comparable displaced and non-displaced workers are in the first place. Appendix 4.F.2 displays our results. Overall, while there exist some differences between displaced and non-displaced workers that could be indicative of workers' families being taken into account in the displacement decision – displaced workers being slightly more likely to be single, and displaced men's female partners being more likely to hold a job than their counterparts – these differences remain limited and are not suggestive of a very strong correlation between observable family-related characteristics and job displacement.

Lastly, we turn back to the comprehensive payroll tax data to get a sense of how likely workers are to leave firms *before* the layoff decision is made. To do so, we plot, over time and separately for each cohort defined by the timing of job displacement (i) the average size of the plant from which workers were displaced, as measured by the number of salaried employees on January, 1st, and (ii) the strength of the outflows from these plants, as measured by the share of workers who no longer work there on December, 31st year T , among those who did on January, 1st year $T - 1$. Figure 4.11 displays our results. Plant size does not diminish gradually over time before the plant-level shock, except perhaps the very last year before the shock. Similarly, outflows appear extremely stable over time before the shock, except one year before the shock. Importantly, this slight anticipation should not affect our results, given that we include in our treated group workers who left the plant either one year after or one year before the plant-level shock.

4.5 Conclusion

In this paper, we quantify the impact of job displacement on workers' family structure, on their spouses' labor supply and lastly on all components that combine into the disposable income as measured at the household level, so as to get a sense of the cost of labor reallocation as paid by workers. Relying on a combination of French administrative registers that span over multiple dimensions of family structure and income, we find that (i) job displacement does not seem to disrupt couples, and does not trigger negative fertility decisions in the short run; (ii) the added worker effect does not seem to prevail in France as displaced workers' partners do not increase their labor supply in response to the shock; (iii) unemployment insurance seems to

provide effective insurance against the short-run income risk that results from job displacement; but (iv) in the long run, unemployment insurance does not seem to matter, as workers are partially insured at best against the long run component, through a bundle that involves self-employment, capital markets and the progressivity of the taxes and transfers system. Importantly, these results depend very little on the gender of the displaced workers.

These empirical results come with a puzzle. First, that households are only partially insured against this risk suggests that they are somewhat credit-constrained or face incomplete markets; otherwise, it would be possible for them to borrow in order to purchase assets of which the returns are negatively correlated with their unemployment risk to get effective insurance against this risk. However, these frictions should make the added worker effect more salient (Lundberg, 1985), which is not observed in the data. An explanation could be that traditional gender norms bias the allocation of time and effort away from the optimal response to these income and employment shocks, but this channel would imply gender-asymmetrical response to one's spouse job loss, which is again not observed in the data. In the short run, unemployment insurance seem to offer good enough insurance to crowd out spousal labor supply, as Cullen and Gruber (2000) would suggest. However, this explanation is no longer viable when it comes to the long run. Solving this puzzle opens avenues for future research.

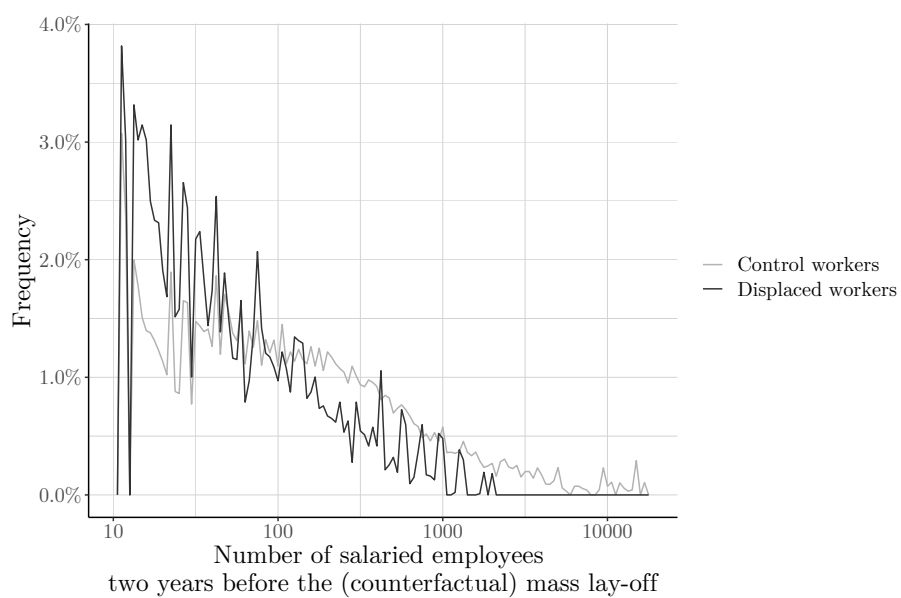
References

- Abadie, A. (2005). Semiparametric Difference-in-Differences Estimators. *The Review of Economic Studies*, 72(1):1–19.
- Abadie, A., Athey, S., Imbens, G. W., and Wooldridge, J. (2017). When should you adjust standard errors for clustering? Working Paper 24003, National Bureau of Economic Research.
- Aghion, P. and Howitt, P. (1994). Growth and unemployment. *The Review of Economic Studies*, 61(3):477–494.
- Akerlof, G. A. and Kranton, R. E. (2000). Economics and Identity. *The Quarterly Journal of Economics*, 115(3):715–753.
- Autor, D., Dorn, D., and Hanson, G. (2019a). When work disappears: Manufacturing decline and the falling marriage market value of young men. *American Economic Review: Insights*, 1(2):161–78.
- Autor, D., Kostøl, A., Mogstad, M., and Setzler, B. (2019b). Disability benefits, consumption insurance, and household labor supply. *American Economic Review*, 109(7):2613–54.
- Becker, G. (1981). *A Treatise on the Family*. Harvard University Press, Cambridge.
- Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2016). Consumption inequality and family labor supply. *American Economic Review*, 106(2):387–435.
- Bozon, M. and Rault, W. (2012). From sexual debut to first union. where do young people in france meet their first partners? *Population*, 67:377–410.
- Brandily, P., Hémet, C., and Malgouyres, C. (2020). Understanding the Reallocation of Displaced Workers to Firms. PSE Working Papers n°2020-82.
- Charles, K. K. and Stephens, Jr., M. (2004). Job displacement, disability, and divorce. *Journal of Labor Economics*, 22(2):489–522.
- Costemalle, V. (2015). Parcours conjugaux et familiaux des hommes et des femmes selon les générations et les milieux sociaux. In *Couples et familles*, Insee Références, pages 63–76. Insee.
- Cullen, J. B. and Gruber, J. (2000). Does unemployment insurance crowd out spousal labor supply? *Journal of Labor Economics*, 18(3):546–572.
- Eliason, M. (2012). Lost jobs, broken marriages. *Journal of Population Economics*, 25(4):1365–1397.
- Gathmann, C., Helm, I., and Schönberg, U. (2020). Spillover Effects of Mass Layoffs. *Journal of the European Economic Association*, 18(1):427–468.
- Halla, M., Schmieder, J., and Weber, A. (2020). Job displacement, family dynamics, and spousal labor supply. *American Economic Journal: Applied Economics*, 12(4):253–87.
- Huttunen, K. and Kellokumpu, J. (2016). The effect of job displacement on couples fertility decisions. *Journal of Labor Economics*, 34(2):403–442.
- Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings losses of displaced workers. *The American Economic Review*, 83(4):685–709.
- Kleven, H., Landais, C., and Sogaard, J. E. (2019). Children and gender inequal-

- ity: Evidence from denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.
- Kleven, H. J., Kreiner, C. T., and Saez, E. (2009). The optimal income taxation of couples. *Econometrica*, 77(2):537–560.
- Lachowska, M., Mas, A., and Woodbury, S. A. (2020). Sources of displaced workers' long-term earnings losses. *American Economic Review*, 110(10):3231–66.
- Lundberg, S. (1985). The added worker effect. *Journal of Labor Economics*, 3(1, Part 1):11–37.
- Rosenfeld, M. J. and Thomas, R. J. (2012). Searching for a mate: The rise of the internet as a social intermediary. *American Sociological Review*, 77(4):523–547.
- Seim, D. (2019). On the incidence and effects of job displacement: Evidence from sweden. *Labour Economics*, 57:131–145.
- Solaz, A., Jalovaara, M., Kreyenfeld, M., Meggiolaro, S., Mortelmans, D., and Pasteeels, I. (2020). Unemployment and separation: Evidence from five european countries. *Journal of Family Research*, 32(1):145–176.
- Stepner, M. (2019). The insurance value of redistributive taxes and transfers.
- Topel, R. H. and Ward, M. P. (1992). Job mobility and the careers of young men. *The Quarterly Journal of Economics*, 107(2):439–479.

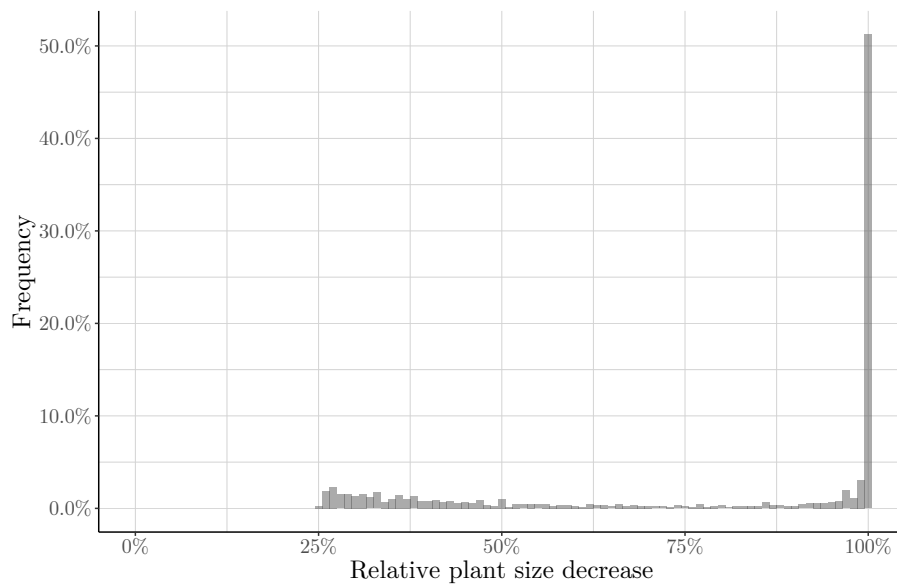
Figures

Figure 4.1 – Plant size before mass layoff



Plant size is measure by the number of salaried employees on January, 1st, one year before the mass layoff.


Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Figure 4.2 – Mass layoff intensity

The mass layoff intensity is 1 minus the number of salaried employees on December, 31st the year of the event, divided by the number of salaried employees on January, 1st the year before.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.


Figure 4.3 – Difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, for those who were before displacement



Rewighted difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFⁱP), EDP sample.


Figure 4.4 – Difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, for those who were not before displacement



Rewighted difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.


Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFⁱP), EDP sample.

Figure 4.5 – Difference-in-difference estimates of the impact of job displacement on the number of children born to workers



Rewighted difference-in-difference estimates of the impact of job displacement on the yearly number of children born to a worker, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications. *Source.* Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

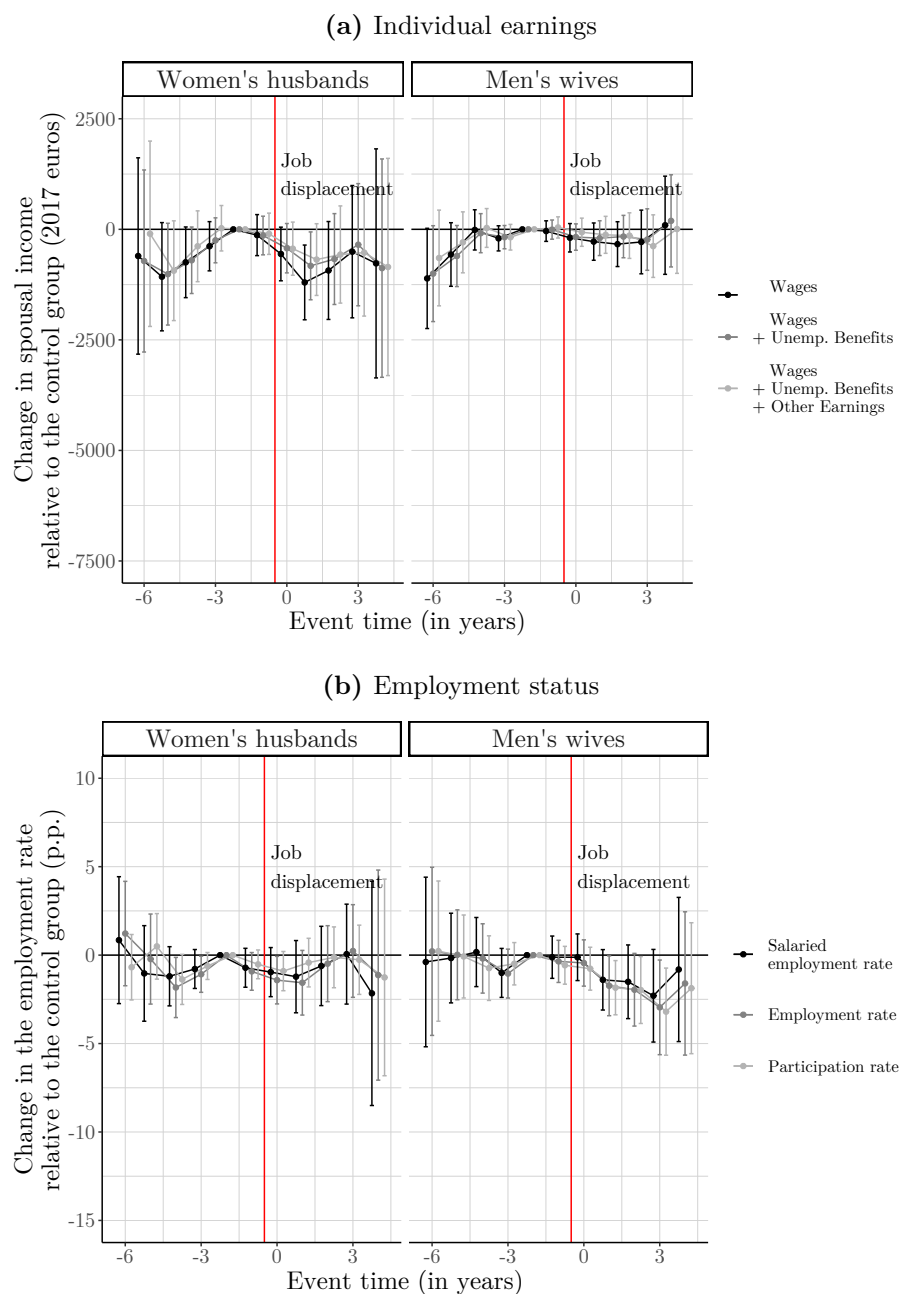
Figure 4.6 – Difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status



Rewighted difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

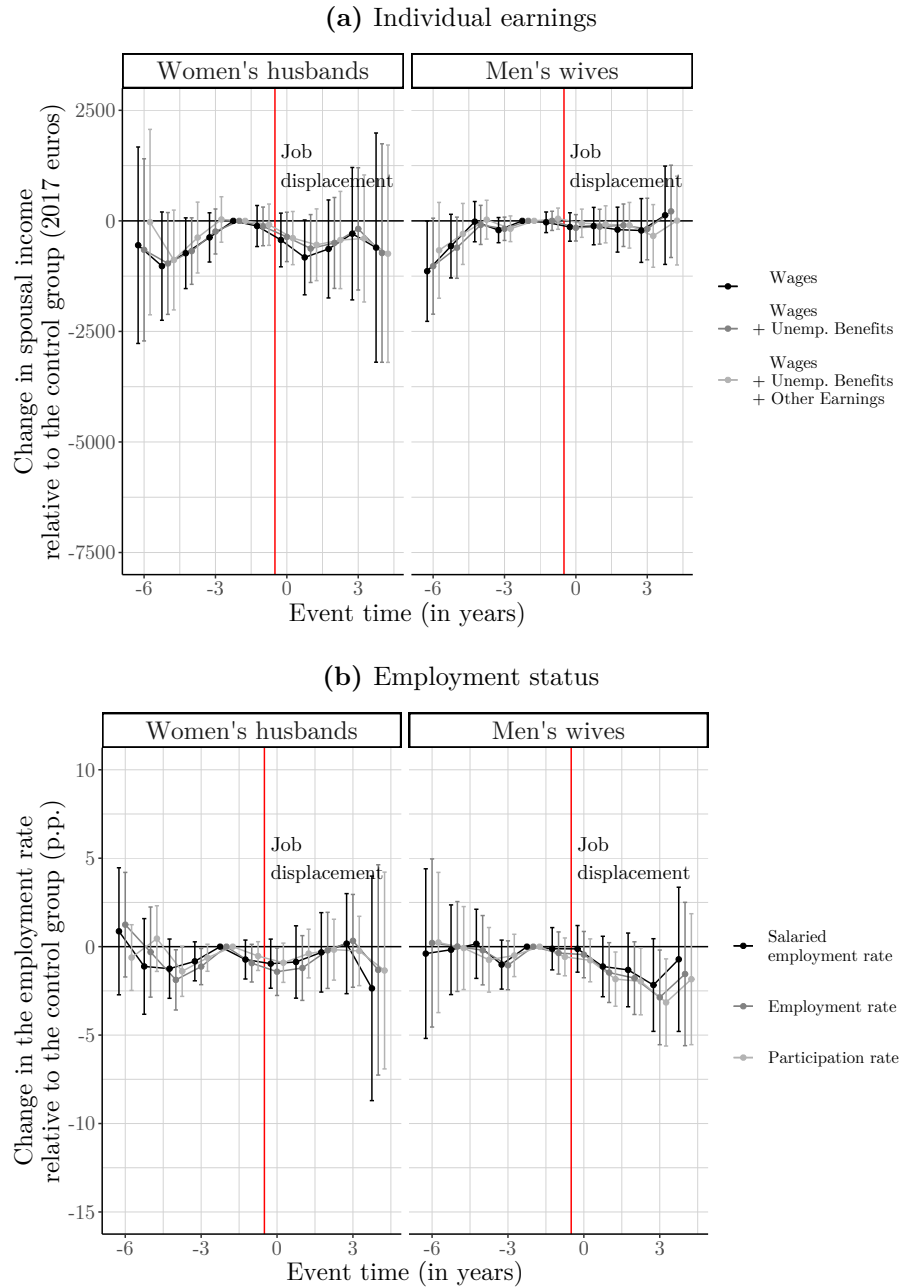
Figure 4.7 – Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status



Rewighted difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

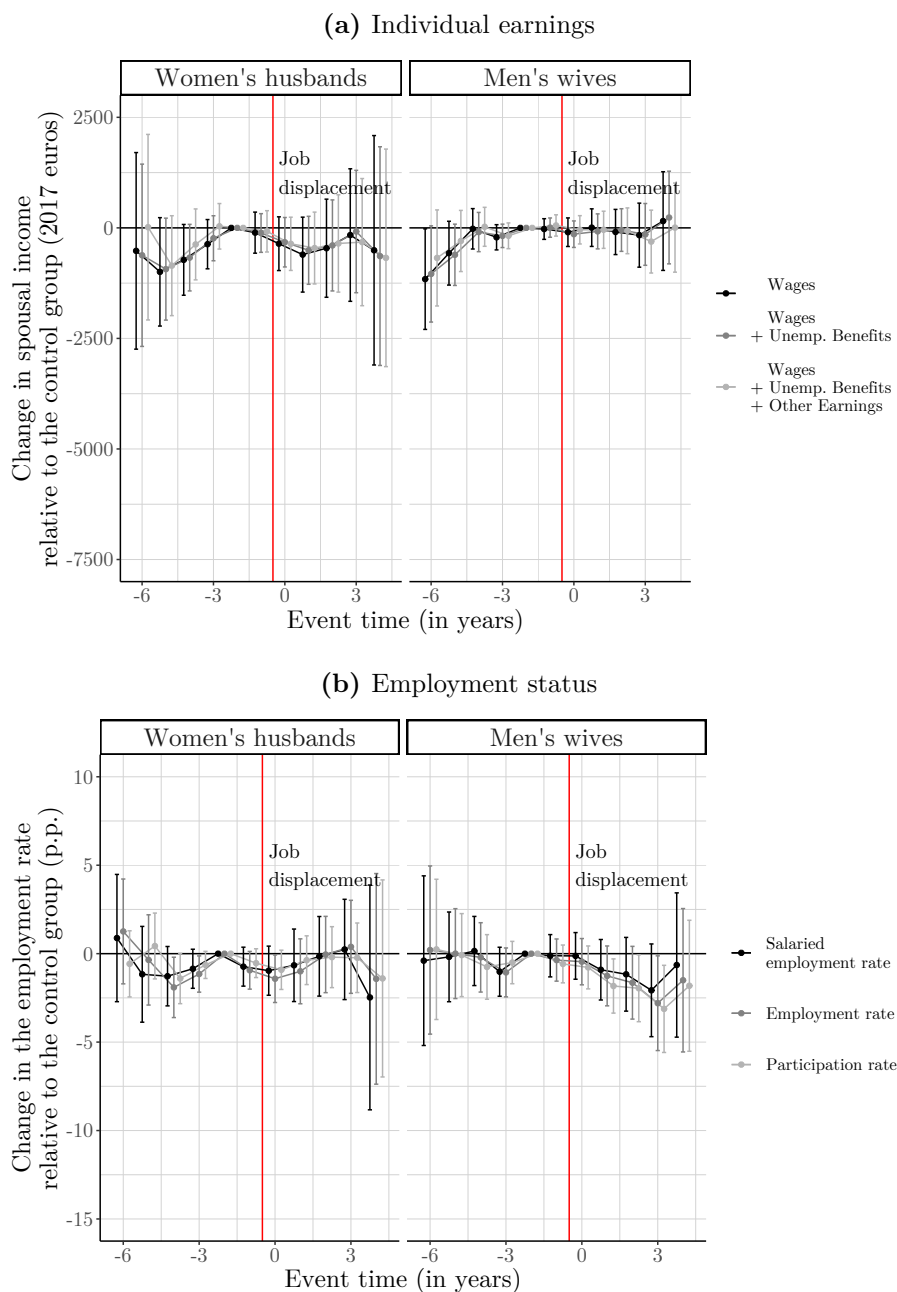
Figure 4.8 – Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the plant level



Rewighted difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.


Figure 4.9 – Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the industry level



Rewighted difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.


Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Figure 4.10 – Difference-in-difference estimates of the impact of job displacement on workers' households' overall income



Rewighted difference-in-difference estimates of the impact of job displacement on workers' households' overall income, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Figure 4.11 – Plant size and outflows over time

Plant size is measured by the number of salaried employees who work at a plant on January, 1st. Outflows correspond to the share of salaried employees who no longer work at the same plant on December, 31st almost two years later.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Tables

Table 4.1 – Summary statistics: occupation and family structure

Gender	Women		Men	
	Control	Displaced	Control	Displaced
<i># Individuals</i>	76,322	3,251	106,348	6,124
<i>a. Age*</i>				
Mean	37.5	37.0	37.8	37.7
St.D.	7.6	7.3	7.6	7.3
<i>b. Occupational shares (in %)**</i>				
Managers and professionals	16.0	19.8	21.2	21.0
Intermediate occupations	21.6	16.8	20.5	17.4
Non-manual workers	48.7	48.2	13.9	15.7
Manual workers	13.7	15.2	44.3	46.0
<i>c. Family structure (in %)**</i>				
Single	11.6	14.6	15.0	14.7
Childless couple	15.2	15.6	14.4	14.1
Single parent	11.5	11.8	5.8	5.7
Couple with children	55.8	50.9	58.5	58.2
Other families	5.8	7.1	6.2	7.3
<i>d. Number of children**</i>				
Mean	1.2	1.1	1.2	1.2
St.D.	1.1	1.2	1.2	1.3
<i>e. Yearly number of childbirths**</i>				
Mean	0.2	0.2	0.2	0.2
St.D.	0.5	0.4	0.5	0.4

* At the time of the (counterfactual) job displacement. ** As observed two years before job displacement. *Source.* Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Table 4.2 – Summary statistics: displaced workers' salaried earnings and labor supply two years before separation (payroll tax data)

Gender	Women		Men	
	Control	Displaced	Control	Displaced
<i>a. Wages (in 2017€)</i>				
Mean	19,900	20,400	26,100	25,200
St.D.	12,100	13,300	14,900	15,600
<i>b. Days worked</i>				
Mean	340	340	340	340
St.D.	70	70	50	60
<i>c. Hours worked</i>				
Mean	1,560	1,580	1,770	1,760
St.D.	500	520	400	500

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Table 4.3 – Summary statistics: displaced workers' earnings and labor supply two years before separation (income tax returns)

Gender	Women		Men	
	Control	Displaced	Control	Displaced
<i>a. Individual earnings (in 2017€)</i>				
Mean	22,000	22,500	27,900	27,100
St.D.	12,300	13,400	15,100	15,700
<i>b. Wages (in 2017€)</i>				
Mean	21,200	21,800	27,300	26,400
St.D.	12,400	13,400	15,200	15,700
<i>c. Unemployment benefits (in 2017€)</i>				
Mean	400	400	300	400
St.D.	1,600	1,800	1,500	1,800
<i>d. Other earnings (in 2017€)</i>				
Mean	400	300	300	300
St.D.	2,200	1,700	2,300	2,000
<i>e. Employment and participation rates (in %)</i>				
Employment	99.6	99.4	99.6	99.6
Participation	99.6	99.5	99.7	99.7

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Table 4.4 – Summary statistics: displaced workers’ spouses’ earnings and labor supply two years before separation (income tax returns)

Gender	Women’s male spouse		Men’s female spouse	
	Control	Displaced	Control	Displaced
<i>a. Individual earnings (in 2017€)</i>				
Mean	30,800	30,900	17,700	16,400
St.D.	19,000	19,600	14,200	14,400
<i>b. Wages (in 2017€)</i>				
Mean	28,000	28,300	16,100	15,100
St.D.	18,700	19,700	13,900	14,200
<i>c. Unemployment benefits (in 2017€)</i>				
Mean	600	600	600	600
St.D.	2,400	2,300	2,100	2,200
<i>d. Other earnings (in 2017€)</i>				
Mean	2,200	2,000	1000	700
St.D.	9,900	8,600	6,200	4,400
<i>e. Employment and participation rates (in %)</i>				
Employment	95.0	95.1	84.0	80.5
Participation	96.2	96.0	85.7	82.6

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Table 4.5 – Summary statistics: displaced workers' household income two years before separation (income tax returns)

Gender	Women		Men	
	Control	Displaced	Control	Displaced
<i>a. Disposable income (in 2017€)</i>				
Mean	46,400	45,900	45,800	45,100
St.D.	24,500	25,700	23,300	24,000
<i>b. Wages (in 2017€)</i>				
Mean	42,000	41,900	41,900	40,700
St.D.	25,600	27,400	24,000	25,100
<i>b. Unemployment benefits (in 2017€)</i>				
Mean	1,100	1,200	1,000	1,200
St.D.	3,100	3,300	2,900	3,200
<i>c. Other earnings (in 2017€)</i>				
Mean	3,700	3,500	3,200	3,400
St.D.	10,600	9,200	9,200	8,200
<i>d. Capital income (in 2017€)</i>				
Mean	4,600	4,400	4,400	4,100
St.D.	8,100	7,800	7,300	7,800
<i>e. Taxes (in 2017€)</i>				
Mean	-7,500	-7,600	-7,300	-7,100
St.D.	7,400	7,900	6,800	7,200
<i>f. Social benefits (in 2017€)</i>				
Mean	2,600	2,500	2,500	2,700
St.D.	3,700	3,600	3,800	4,000
<i>g. Position in the equivalent income distribution (in %)</i>				
First quintile	9.1	9.8	9.5	12.7
Second quintile	20.6	21.7	19.6	21.3
Third quintile	24.1	22.1	24.6	23.0
Fourth quintile	24.7	23.2	25.8	23.1
Fifth quintile	21.4	23.2	20.5	19.9

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

4.A Institutional Background

In France, mass layoffs for economic reasons (i.e. layoff of at least two employees) are tightly regulated.

Duration. A single mass layoff is defined as the dismissal of two or more employees over 30 days. The time lapse between the notification of the layoff and the termination of the contract depends on the tenure of the worker: one month for workers with a tenure between 6 months and two years, two months for those with a tenure of more than two years.¹⁴ A single mass layoff can spread over three months and even more if, before being notified of the layoff, the worker had scheduled paid vacations during the period of notice.

Order of dismissals If a collective agreement applicable to the firm specifies the criteria determining the order of dismissals, then they are binding on the employer. Otherwise, the employer must define them, after consulting the *Comité Social et Économique* (CSE). In this case, the employer is bound take into account all of the following elements: (i) employees' family responsibilities, in particular single parenthood; (ii) employees tenure in the plant or firm; (iii) any situation making re-employment especially difficult, in particular that of elderly or disabled employees; and lastly (iv) skills assessed by occupation (e.g. speed and versatility in carrying out tasks). Other criteria may be added to this list, and the employer may give preference to one of these criteria, but this is only possible provided that other criteria are taken into account. Additionally, these criteria have to be taken into account even when the stake is an individual dismissal for economic reasons.

Lastly, anti-discrimination laws prohibit that certain criteria be used for the decisions, and it is also not possible to dismiss an employee solely because he or she works part-time in preference to a full-time employee performing the same type of tasks in the company.

Employment protection program. Firms with at least 50 employees and firing at least 10 employees over 30 days have to implement a *Plan de Sauvegarde de l'Emploi* (PSE). This employment protection program includes several measures to

¹⁴Below 6 months, the length of the notice is either set by collective agreements or depends on practices in the company or the industry.

support the reemployment of dismissed employees (internal or external redeployment, formations,...) When the firm implements a PSE, severance payments are exempted from income taxes.

Compensation. First, upon termination of a permanent contract¹⁵ for economic reasons (mass layoff or firm dissolution), workers with tenure 8 months or more are entitled to severance payments. These severance payments are bound to a minimum which depends on their gross monthly reference wage¹⁶ w according to the following formula:

$$\frac{1}{4}w \min \left\{ \frac{M}{12}, 10 \right\} + \frac{1}{3}w \max \left\{ \frac{M}{12} - 10, 0 \right\}$$

Firm- or industry-level agreements may set higher minima. These severance payments are the same for collective contractual termination as they are for standard layoffs. On top of these payments, workers may be eligible to compensations for periods of notice and paid vacations they could not take advantage of before termination of their contract. These periods are compensated at a wage rate and these payments are considered as wages by the fiscal administration.

Legal severance payments as well as those set by collective agreements are fully exempted from income taxes and partly exempted from social contributions. Supra-legal payments are partly exempted from income taxes (except in the case of a PSE where they are fully exempted) and are not exempted from social contributions.

Firms below 1,000 employees (or firms going into receivership) have to offer a *Contrat de Sécurisation Professionnelle (CSP)* to dismissed workers.¹⁷ Upon their acceptance, these workers are entitled to a one-year special unemployment compensation called the *Allocation de Sécurisation Professionnelle (ASP)*. This compensation is the same as regular unemployment insurance benefits for workers with tenure less than one year, i.e. 57% of the previous wage in the most standard case. For workers with tenure over one year, this compensation amounts to 75% of their previous wage. Should she fail to find a new job during this time span, an ASP receiver is entitled to regular UI benefits for a potential duration equal to her initial entitlement duration upon job loss, minus one year.

¹⁵ *Contrat à Durée Indéterminée (CDI)*

¹⁶ This reference wage is defined as the higher value between average gross wage computed either over the past year or over the past three months.

¹⁷ There is no CSP when the employer and the employee sign an agreement for a contractual termination.

Firms with over 1,000 employees have to offer a *Congé de reclassement*, which provides training and job search support. The length of this leave is between 4 and 12 months (24 months in case of a professional retraining). The employer pays the full wage during the period of notice and then at least 65% of the reference wage, with a lower limit equal to 85% of the French minimum wage (SMIC). This compensation is taxed as a labor income.

Finally, in order to cut labor costs, firms may also arrange early terminations of fixed-term contracts. At the end of a fixed-term contract, employees receive an allowance equal to at least 10% of the total gross income they earned during this contract. However, during a mass layoff, fixed-term contracts may be terminated before the end of the contract. When this decision is initiated by the employer, the employee receives severance payments equal to the sum of earnings she would have perceived until the end of this contract. During the trial period, the fixed-term contract can be terminated without compensation.

4.B Discouraged worker effect

Our approach has so far dealt with the identification of shocks on the employment status of one worker on her and her spouse's outcomes. As explained in Subsection 4.3.2, the issue with this framework is that the spouse is likely directly affected by employment shocks that are positively correlated with the shocks of which we aim at identifying the consequences. In this Appendix, we develop a very simple model of couples exposed to correlated employment shocks to get a plausible approximation of the size of the resulting bias.

We consider a population of heterosexual couples, indexed by i , which are composed of a female and a male partner $f(i)$ and $m(i)$. Each partner (i) has outcomes $Y_{f(i)}$ and $Y_{m(i)}$ respectively; and (ii) is exposed to employment shocks that are represented by dummy variables $D_{f(i)}$ and $D_{m(i)}$.

We frame the model in terms of potential outcomes. By contrast with our initial framework, in which potential outcomes are made to depend only on the observed partner's employment shocks, here these potential reactions, denoted as $Y_{f(i)}(d_{f(i)}, d_{m(i)})$ and $Y_{m(i)}(d_{f(i)}, d_{m(i)})$ may depend jointly on both partners' employment shocks. In order to lighten the exposition of the model, we make the following simplifying assumption:

Assumption 4.B.1 (Exogeneity). *Couples' potential outcomes are mean independent of couples employment shocks: for all g in $\{f, m\}$, for all $(d_{f(i)}, d_{m(i)})$ and $(d'_{f(i)}, d'_{m(i)})$ in $\{0, 1\}^2$:*

$$\mathbb{E}[Y_{g(i)}(d_{f(i)}, d_{m(i)}) \mid D_{f(i)} = d'_{f(i)}, D_{m(i)} = d'_{m(i)}] = \mathbb{E}[Y_{g(i)}(d_{f(i)}, d_{m(i)})] \quad (4.8)$$

This assumption states that employment shocks can be treated as though they are random and uncorrelated with couples potential reaction to them. With respect to our conditional difference-in-difference design, this assumption is a massive simplification; it basically holds (i) after conditioning on observables, and differencing with respect to the pre-shock period outcomes; and (ii) assuming that average treatment effects are the same for treated and control groups. The differencing and conditioning remains implicit in our one-period model so as to keep the notations handy.

For any gender g in $\{f, m\}$ and any two vectors $(d_{f(i)}, d_{m(i)})$ and $(d'_{f(i)}, d'_{m(i)})$,

we define the average joint treatment effect:

$$\Delta_{(d_{f(i)}, d_{m(i)}) \rightarrow (d'_{f(i)}, d'_{m(i)})}^g = \mathbb{E}[Y_{g(i)}(d'_{f(i)}, d'_{m(i)}) - Y_{g(i)}(d_{f(i)}, d_{m(i)})] \quad (4.9)$$

This quantity represents the average effect, for partners of gender g , of moving from the $(d_{f(i)}, d_{m(i)})$ situation to the $(d'_{f(i)}, d'_{m(i)})$ situation.

For g and g' in $\{f, m\}$, we also define the average partial treatment effect:

$$\Delta^{(g, g')} = \mathbb{E}[Y_g(i)(d_{g'(i)} = 1, D_{-g'(i)}) - Y_g(i)(d_{g'(i)} = 0, D_{-g'(i)})] \quad (4.10)$$

This quantity represents the average effects for partners of gender g 's outcomes of a shock on partners of gender g' 's employment relationship, *without changes in the shock that affects the other partner's status*.

Armed with these definitions and our exogeneity assumption, we are now able to show that the within couple correlation in employment shocks does indeed generate bias in our estimates, even when employment shocks are exogenous.

Proposition 4.B.1. *Under Assumption 4.B.1, the difference in expected outcomes the g' -treated and the g' -control group does not identify the partial treatment effect, unless (i) the other partner's employment shocks have no effect on the outcome, or (ii) there is no within couple correlation in employment shocks.*

Proof. Without loss of generality, we assume $g' = f$, i.e. that we are investigating

the consequences of a shock on female employment:

$$\begin{aligned}
& \mathbb{E}[Y_{g(i)} \mid D_{f(i)} = 1] - \mathbb{E}[Y_{g(i)} \mid D_{f(i)} = 0] \\
= & \mathbb{E}[Y_{g(i)} \mid D_{f(i)} = 1, D_{m(i)} = 1] \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 1) \\
& + \mathbb{E}[Y_{g(i)} \mid D_{f(i)} = 1, D_{m(i)} = 0] \{1 - \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 1)\} \\
& - \mathbb{E}[Y_{g(i)} \mid D_{f(i)} = 1, D_{m(i)} = 0] \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 0) \\
& - \mathbb{E}[Y_{g(i)} \mid D_{f(i)} = 0, D_{m(i)} = 0] \{1 - \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 0)\} \\
\stackrel{A4.B.1}{=} & \mathbb{E}[Y_{g(i)}(1, 1)] \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 1) \\
& + \mathbb{E}[Y_{g(i)}(1, 0)] \{1 - \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 1)\} \\
& - \mathbb{E}[Y_{g(i)}(0, 1)] \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 0) \\
& - \mathbb{E}[Y_{g(i)}(0, 0)] \{1 - \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 0)\} \\
= & \mathbb{E}[Y_{g(i)}(1, 1) - Y_{g(i)}(0, 1)] \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 1) \\
& + \mathbb{E}[Y_{g(i)}(1, 0) - Y_{g(i)}(0, 0)] \{1 - \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 1)\} \\
& + \mathbb{E}[Y_{g(i)}(0, 1) - Y_{g(i)}(0, 0)] \{\mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 1) - \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 0)\} \\
= & \Delta^{(g,f)} \\
& + \Delta_{(0,0) \rightarrow (0,1)}^g \{\mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 1) - \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 0)\} \tag{4.11}
\end{aligned}$$

where the first equality follows from the law of iterated expectations, the third one from the linearity of the expectation operator and the last one from the law of iterated expectations and the definition of the average partial and joint treatment effects.

If shocks on the employment status of men have no effect on $Y_{g(i)}$, then $\Delta_{(0,0) \rightarrow (0,1)}^g = 0$. If there is no within-couple correlation in employment shocks, then $\mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 1) = \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 0)$. Thus in either cases, $\mathbb{E}[Y_{g(i)} \mid D_{f(i)} = 1] - \mathbb{E}[Y_{g(i)} \mid D_{f(i)} = 0] = \Delta^{(g,f)}$. Otherwise, the bias corresponds to the remaining term in Equation 4.11. \square

The remainder of this Appendix is devoted to the general case, that is the one in which employment shocks are assumed to be correlated within couple. Proposition 4.B.1 states that our naive approach that simply compares couples of which partners of gender g are affected by a shock to couples of which the same partners are not affected by a shock is biased due to the fact that (i) the other partners in the treated couples are more likely to be affected by an employment shock (assuming the

correlation is positive), and (ii) this additional shock changes couples' labor supply decisions. Our aim is to offer a reasonable approximation for the size of the bias.

In order to do so, we make several simplifying assumptions:

Assumption 4.B.2 (Direct effects). *Direct effects are much larger than indirect effects, i.e.:*

$$\begin{cases} \Delta_{(0,0) \rightarrow (0,1)}^m, \Delta_{(1,0) \rightarrow (1,1)}^m \gg \Delta_{(0,0) \rightarrow (1,0)}^m \\ \Delta_{(0,0) \rightarrow (1,0)}^f, \Delta_{(0,1) \rightarrow (1,1)}^f \gg \Delta_{(0,0) \rightarrow (0,1)}^f \end{cases} \quad (4.12)$$

Assumption 4.B.3 (Limited correlation). *The within couple correlation in employment shocks is limited, i.e. for all g in $\{m, f\}$:*

$$\mathbb{P}(D_{-g(i)} = 1 \mid D_{g(i)} = 1) - \mathbb{P}(D_{-g(i)} = 1 \mid D_{g(i)} = 0) \ll 1 \quad (4.13)$$

Assumption 4.B.2 implies that a worker's labor outcomes reacts much more to whether she is affected by an employment shock than to whether her partner is. Assumption 4.B.3 states that most couples evolve in somewhat isolated labor markets, so that most of the time a worker being hit by a shock does not imply that her spouse is too.

As implied by Assumption 4.B.2, we make the distinction between direct effects, i.e. the impact of one's own employment shock on one's own labor outcomes, and indirect effects, i.e. the impact of a shock on one's spouse employment relationship. When considering direct shocks under these new assumptions, the size of the bias as it appears in Equation 4.11 is small before $\Delta^{g,g}$. As a result, our approach offers a good approximation of the direct effect of employment shocks.

The same does not hold when it comes to indirect effects: indeed, this time the bias results from the multiplication of large direct effects by the within couple correlation, which can no longer be assumed to be small before $\Delta^{g,-g}$. To get a plausible approximation for the size of the bias, we introduce additional assumptions:

Assumption 4.B.4 (Limited treatment effect heterogeneity). *Direct effects depend very little on whether the spouse is affected by a shock:*

$$\begin{cases} \Delta_{(0,0) \rightarrow (0,1)}^m \simeq \Delta_{(1,0) \rightarrow (1,1)}^m \\ \Delta_{(0,0) \rightarrow (1,0)}^f \simeq \Delta_{(0,1) \rightarrow (1,1)}^f \end{cases} \quad (4.14)$$

Assumption 4.B.4 implies in particular that our estimates of the aggregate direct effect $\Delta^{(g,g)}$ is a reasonable approximation for the first component of the bias as it

appears in Equation 4.11.

To approximate the second component of the bias, we now assume that we are provided with a certain partition of couples that predicts perfectly the within couple correlation in employment shocks. Specifically, we assume that this partition corresponds to the distinction between couple in which both partners work in the same firm (detailed industry), as opposed to partners who work in different firms (detailed industry).

Assumption 4.B.5 (Predicted correlation). *There exists an observable subset of couples \mathcal{J} such that (i) if i belongs to \mathcal{J} , then spouses' employment shocks are perfectly correlated: $D_{f(i)} = D_{m(i)}$; and (ii) otherwise spouses' employment shocks are independent: $D_{f(i)} \perp\!\!\!\perp D_{m(i)}$.*

Assumption 4.B.6 (Limited treatment probability heterogeneity). *The probability than one spouse is affected by an employment shock depends very little on whether shocks are correlated or not between spouses: for all g in $\{f, m\}$*

$$\mathbb{P}(D_{g(i)} = 1 \mid i \in \mathcal{J}) \simeq \mathbb{P}(D_{g(i)} = 1 \mid i \notin \mathcal{J}) \quad (4.15)$$

These assumptions provide us with approximations of the second component of the bias.

Proposition 4.B.2 (Bias approximation). *Under Assumptions 4.B.1 to 4.B.6, the bias that results from correlated employment shocks between spouses in the estimation of indirect effects is approximately equal to the direct effect multiplied by the share of couples that belong to \mathcal{J} .*

Proof. We assume without loss of generality that we focus on the impact of a female employment shock on men's labor outcomes. That $\Delta_{(0,0) \rightarrow (0,1)}^m$ is well approximated by $\Delta^{(m,m)}$ is a direct consequence of Assumption 4.B.4 combined with the law of iterated expectations. This gives us a good estimate of the first term of the bias.

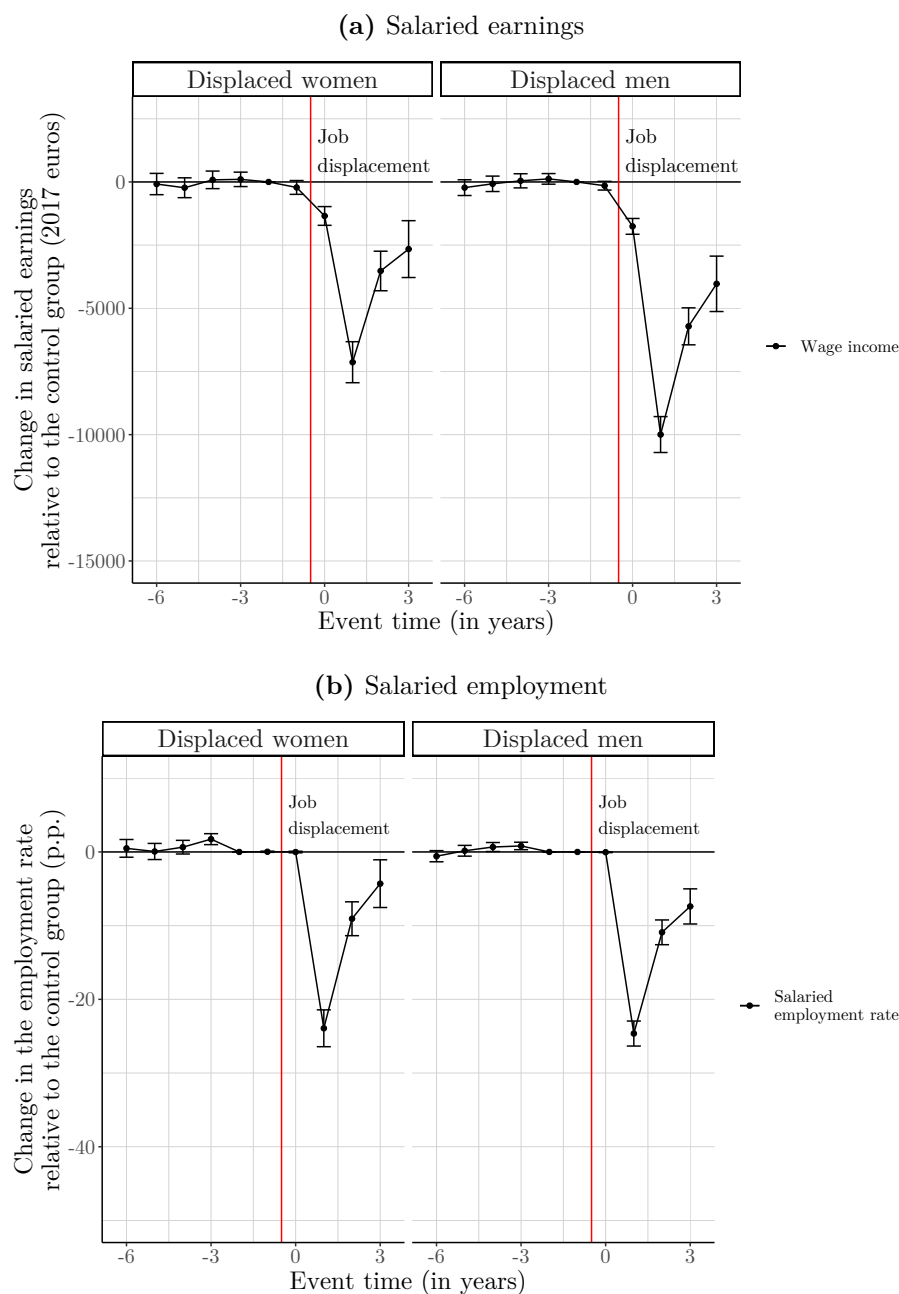
The second component of the bias writes:

$$\begin{aligned}
& \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 1) - \mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 0) \\
= & \mathbb{P}(i \in \mathcal{J} \mid D_{f(i)} = 1)\mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 1, i \in \mathcal{J}) \\
& + \{1 - \mathbb{P}(i \in \mathcal{J} \mid D_{f(i)} = 1)\}\mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 1, i \notin \mathcal{J}) \\
& - \mathbb{P}(i \in \mathcal{J} \mid D_{f(i)} = 0)\mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 0, i \in \mathcal{J}) \\
& - \{1 - \mathbb{P}(i \in \mathcal{J} \mid D_{f(i)} = 0)\}\mathbb{P}(D_{m(i)} = 1 \mid D_{f(i)} = 0, i \notin \mathcal{J}) \\
\stackrel{A4.B.5}{=} & \mathbb{P}(i \in \mathcal{J} \mid D_{f(i)} = 1) \\
& + \{1 - \mathbb{P}(i \in \mathcal{J} \mid D_{f(i)} = 1)\}\mathbb{P}(D_{m(i)} = 1 \mid i \notin \mathcal{J}) \\
& - \{1 - \mathbb{P}(i \in \mathcal{J} \mid D_{f(i)} = 0)\}\mathbb{P}(D_{m(i)} = 1 \mid i \notin \mathcal{J}) \\
\stackrel{A4.B.6}{=} & \mathbb{P}(i \in \mathcal{J}) \tag{4.16}
\end{aligned}$$

where the first equality simply follows from the law of iterated expectations. \square

4.C Earnings and employment effects in the DADS panel

Figure 4.C.1 – Difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status




Rewighted difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

4.D Individual-level income loss for displaced workers with spouses

Figure 4.D.1 – Difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status (workers in a cohabiting heterosexual relationship)




Rewighted difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

4.E Sample selection

4.E.1 Impact of job displacement on sample inclusion

Figure 4.E.1 – Difference-in-difference estimates of the impact of job displacement on workers' presence in the income tax returns data

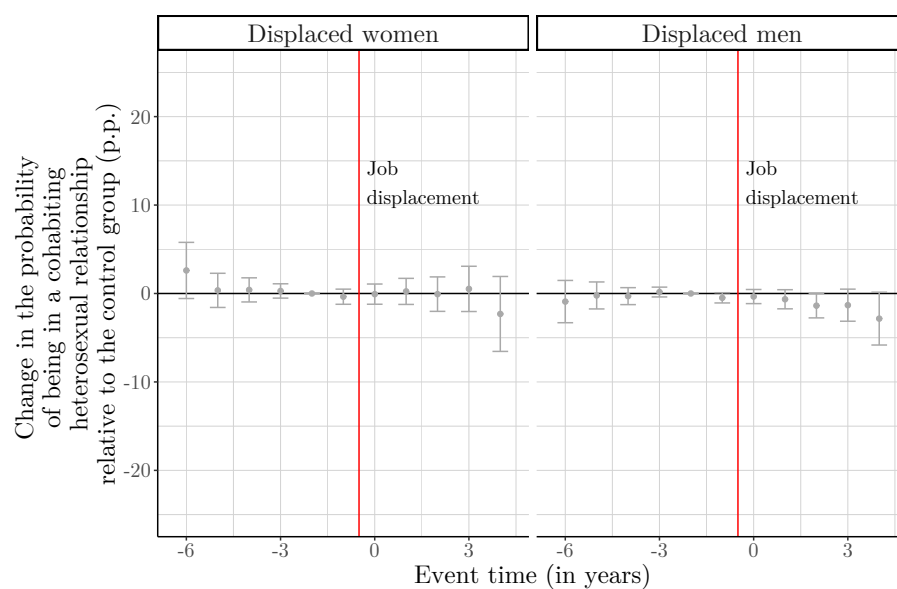


Rewighted difference-in-difference estimates of the impact of job displacement on workers' presence in income tax returns, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFⁱP), EDP sample.

4.E.2 Replication on the balanced panel

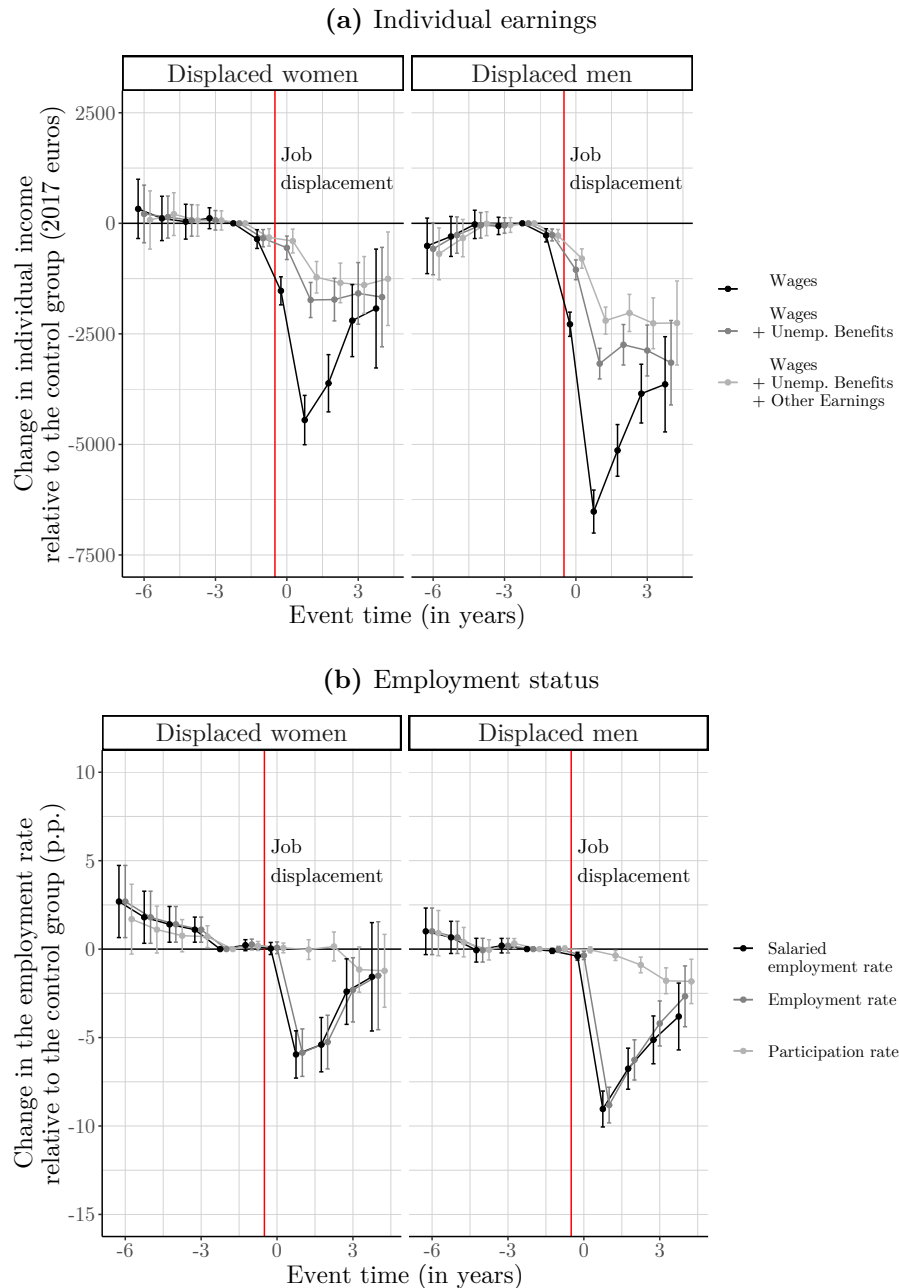
Figure 4.E.2 – Difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship (balanced panel)



Rewighted difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.


Figure 4.E.3 – Difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status (balanced panel)



Rewighted difference-in-difference estimates of the impact of job displacement on workers' own individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

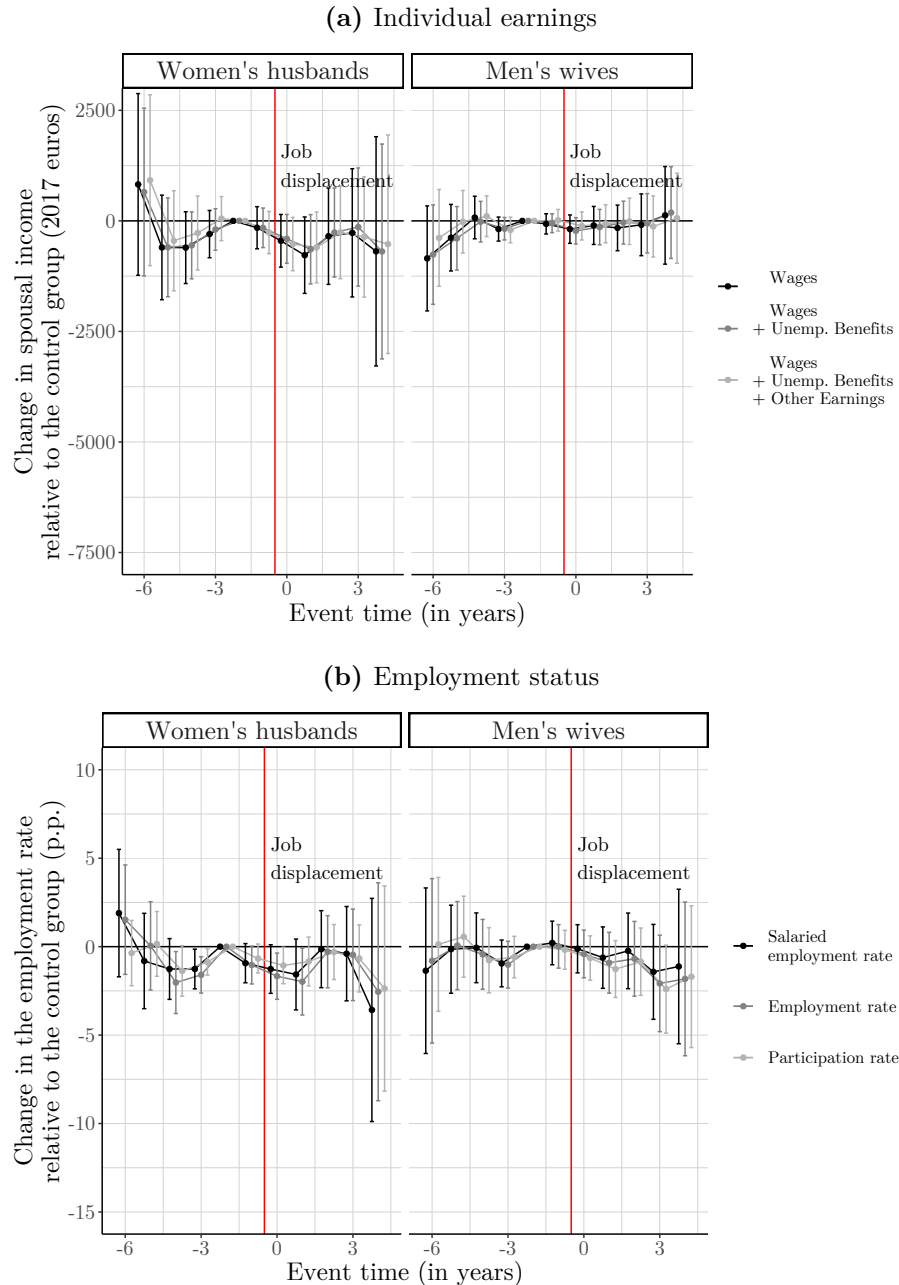
Figure 4.E.4 – Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status (balanced panel)



Rewighted difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

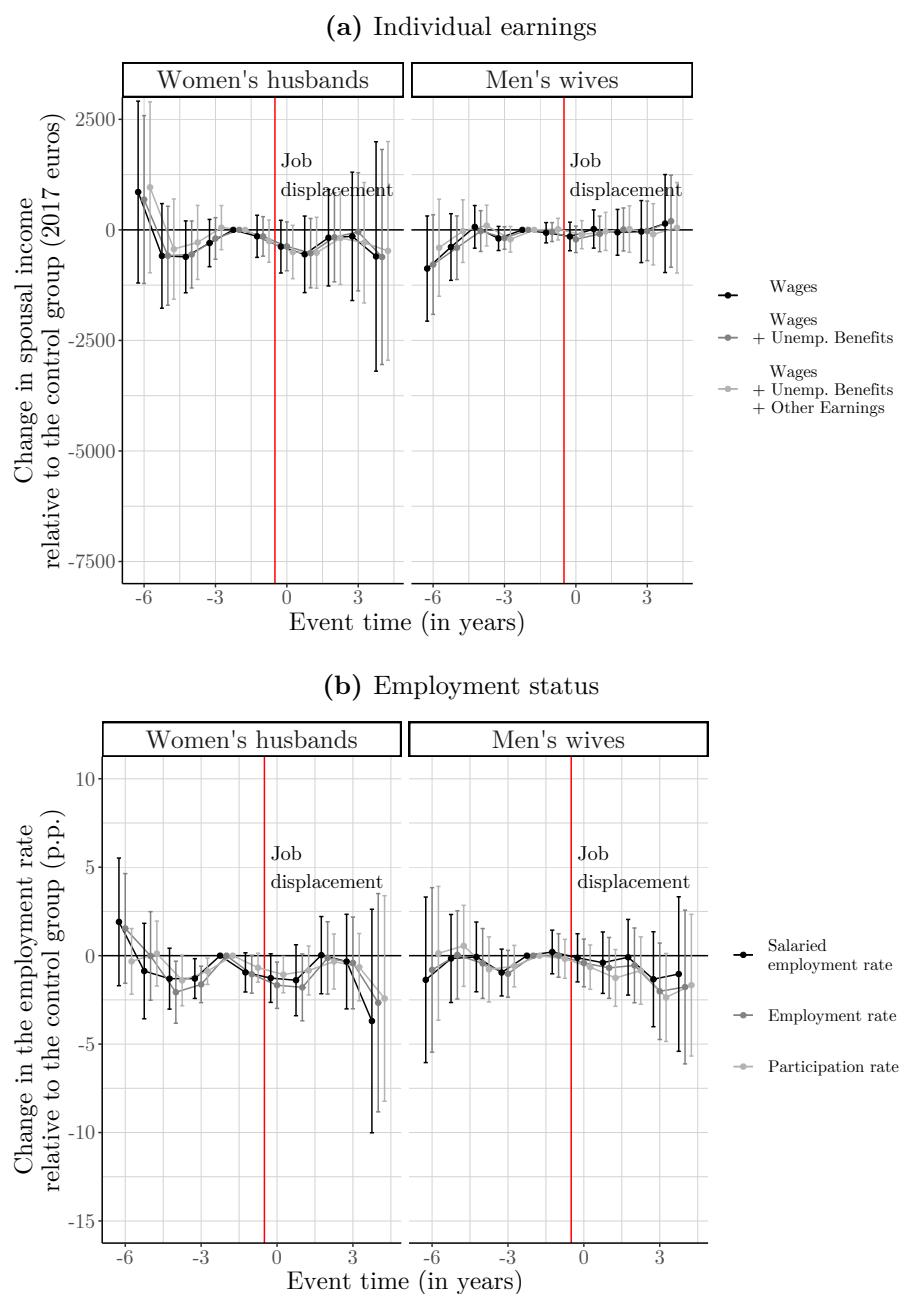
Figure 4.E.5 – Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the plant level (balanced panel)



Rewighted difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.


Figure 4.E.6 – Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the industry level (balanced panel)



Rewighted difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Figure 4.E.7 – Difference-in-difference estimates of the impact of job displacement on workers' households' overall income (balanced panel)

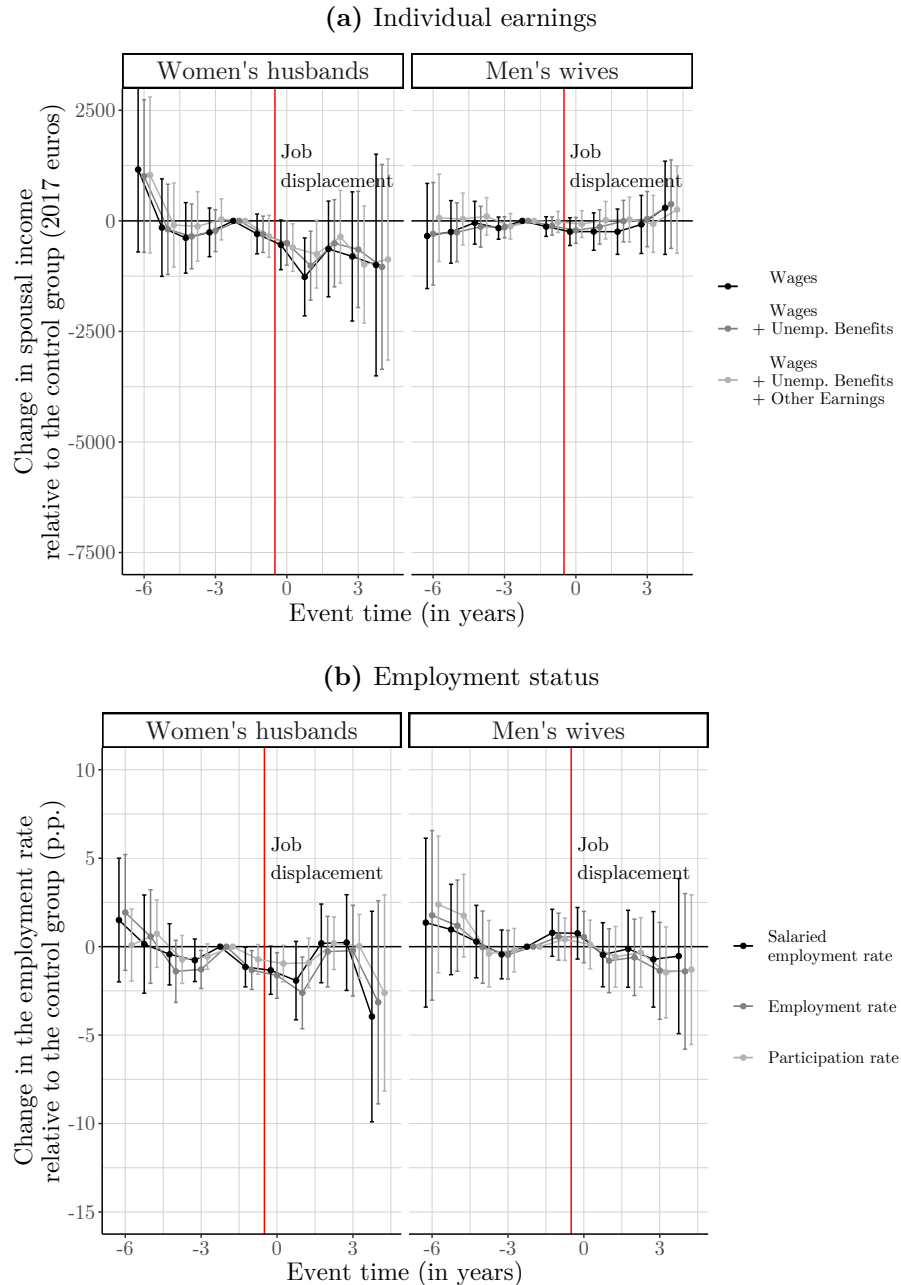


Rewighted difference-in-difference estimates of the impact of job displacement on workers' households' overall income, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

4.E.3 Replication on the balanced panel of spouses

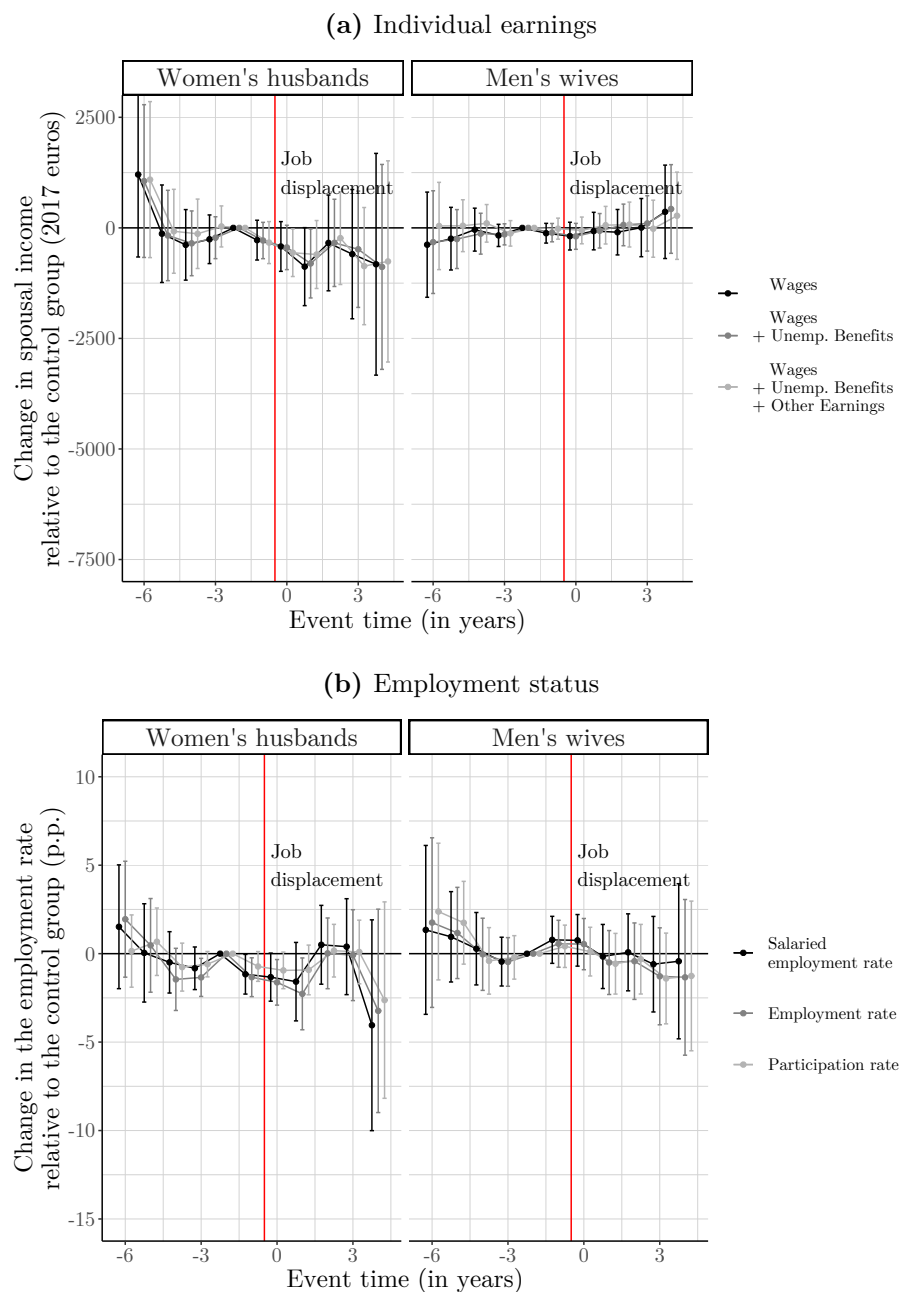
Figure 4.E.8 – Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status (balanced panel of spouses)



Rewighted difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

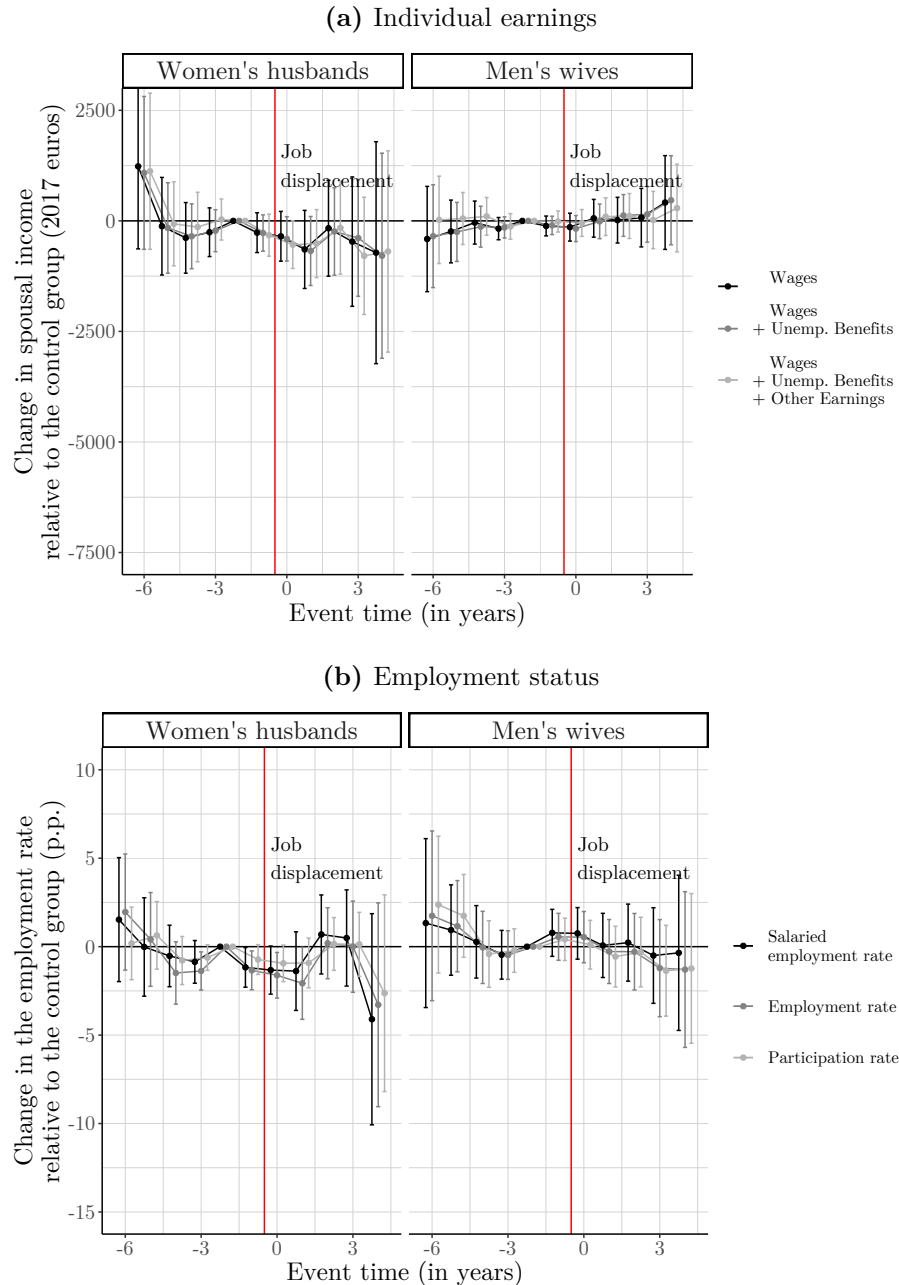
Figure 4.E.9 – Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the plant level (balanced panel)



Rewighted difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Figure 4.E.10 – Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the industry level (balanced panel)




Rewighted difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

4.F Selection into job displacement

4.F.1 Restriction to plant-closure events


Figure 4.F.1 – Difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, for those who were before displacement (plant-closure events only)



Rewighted difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.


Figure 4.F.2 – Difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, for those who were not before displacement (plant-closure events only)



Rewighted difference-in-difference estimates of the impact of job displacement on the probability of being in a heterosexual cohabiting relationship, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

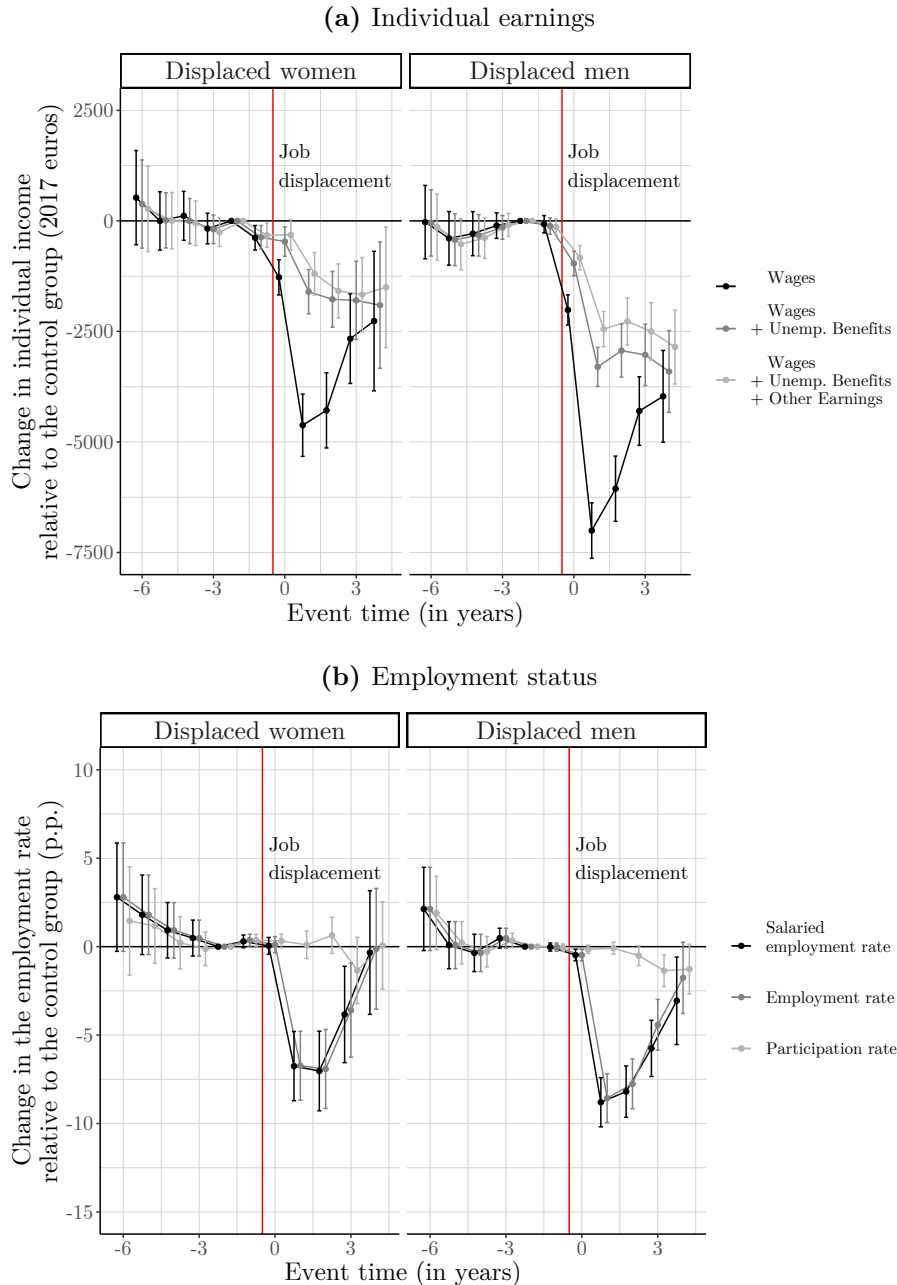
Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFⁱP), EDP sample.

Figure 4.F.3 – Difference-in-difference estimates of the impact of job displacement on the number of children born to workers (plant-closure events only)



Rewighted difference-in-difference estimates of the impact of job displacement on the yearly number of children born to a worker, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications. *Source.* Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

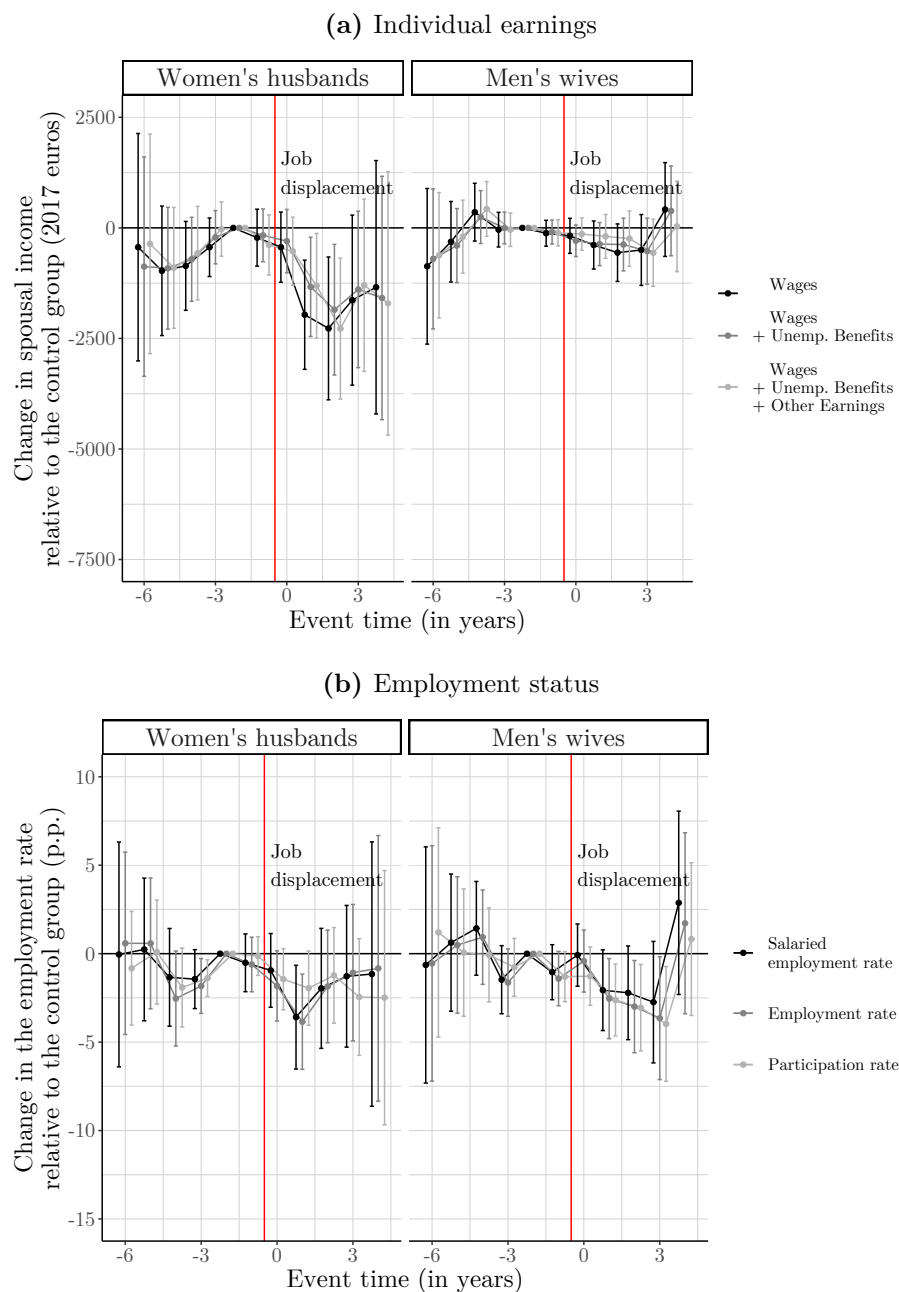
Figure 4.F.4 – Difference-in-difference estimates of the impact of job displacement on workers’ own individual earnings and employment status (plant-closure events only)



Rewighted difference-in-difference estimates of the impact of job displacement on workers’ own individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

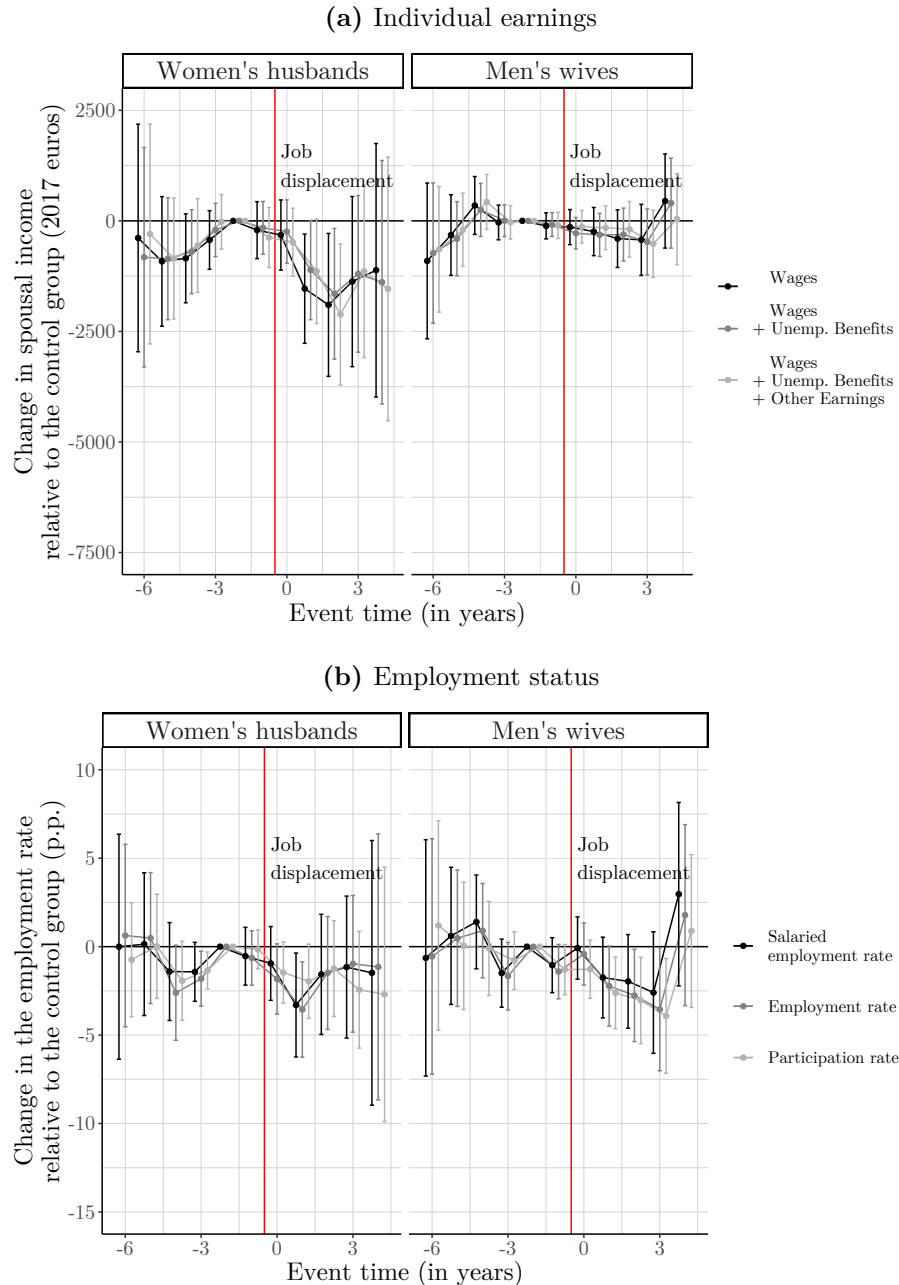
Figure 4.F.5 – Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status (plant-closure events only)



Rewighted difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.


Figure 4.F.6 – Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the plant level (plant-closure events only)



Rewighted difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.


Figure 4.F.7 – Difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, with the discouraged worker effect correction at the industry level (plant-closure events only)



Rewighted difference-in-difference estimates of the impact of job displacement on workers' spouses' individual earnings and employment status, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Figure 4.F.8 – Difference-in-difference estimates of the impact of job displacement on workers' households' overall income (plant-closure events only)



Rewighted difference-in-difference estimates of the impact of job displacement on workers' households' overall income, by gender of the displaced worker and time since job displacement. Standard errors are clustered at the plant level and estimated by bootstrap with 200 replications.

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

4.F.2 Comparison with non-displaced workers of downsized plants

Table 4.F.1 – Summary statistics: occupation and family structure

Gender	Women		Men	
	Non-displaced	Displaced	Non-displaced	Displaced
<i># Individuals</i>	578	1,702	1,068	3,001
<i>a. Age</i> *				
Mean	36.6	36.3	37.5	37.4
St.D.	7.0	7.2	7.1	7.2
<i>b. Occupational shares (in %)**</i>				
Managers and professionals	18.5	23.8	16.7	25.1
Intermediate occupations	15.4	17.3	17.7	17.8
Non-manual workers	52.2	45.9	15.3	17.1
Manual workers	13.8	12.9	50.4	40.1
<i>c. Family structure (in %)**</i>				
Single	12.5	16.6	14.0	15.8
Childless couple	18.2	16.0	13.9	14.7
Single parent	8.5	11.9	5.9	5.7
Couple with children	55.5	48.1	58.9	57.1
Other families	5.4	7.4	7.4	6.7
<i>d. Number of children**</i>				
Mean	1.2	1.0	1.2	1.2
St.D.	1.1	1.1	1.2	1.2
<i>e. Yearly number of childbirths**</i>				
Mean	0.2	0.2	0.2	0.2
St.D.	0.4	0.5	0.5	0.5

* At the time of the (counterfactual) job displacement. ** As observed two years before job displacement. *Source.* Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Table 4.F.2 – Summary statistics: displaced workers' salaried earnings and labor supply two years before separation (payroll tax data)

Gender	Women		Men	
	Non-displaced	Displaced	Non-displaced	Displaced
<i>a. Wages (in 2017€)</i>				
Mean	19,900	20,400	26,100	25,200
St.D.	12,100	13,300	14,900	15,600
<i>b. Days worked</i>				
Mean	340	340	340	340
St.D.	70	70	50	60
<i>c. Hours worked</i>				
Mean	1,560	1,580	1,770	1,760
St.D.	500	520	400	500

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Table 4.F.3 – Summary statistics: displaced workers' earnings and labor supply two years before separation (income tax returns)

Gender	Women		Men	
	Non-displaced	Displaced	Non-displaced	Displaced
<i>a. Individual earnings (in 2017€)</i>				
Mean	23,500	23,600	25,900	28,700
St.D.	14,600	14,700	14,400	17,400
<i>b. Wages (in 2017€)</i>				
Mean	22,400	22,700	25,000	28,000
St.D.	14,300	14,700	14,00	17,500
<i>c. Unemployment benefits (in 2017€)</i>				
Mean	600	500	600	500
St.D.	2,100	1,900	2,200	1,900
<i>d. Other earnings (in 2017€)</i>				
Mean	400	300	200	300
St.D.	3,500	1,900	2,000	1,900
<i>e. Employment and participation rates (in %)</i>				
Employment	98.3	99.5	98.7	99.6
Participation	98.3	99.6	99.4	99.6

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Table 4.F.4 – Summary statistics: displaced workers' spouses' earnings and labor supply two years before separation (income tax returns)

Gender	Women's male spouse		Men's female spouse	
	Non-displaced	Displaced	Non-displaced	Displaced
<i>a. Individual earnings (in 2017€)</i>				
Mean	33,000	32,500	15,200	17,800
St.D.	21,000	19,900	12,700	15,200
<i>b. Wages (in 2017€)</i>				
Mean	29,100	30,300	13,700	16,300
St.D.	21,400	20,300	12,700	15,200
<i>c. Unemployment benefits (in 2017€)</i>				
Mean	600	500	600	600
St.D.	2,400	2,200	2,000	2,200
<i>d. Other earnings (in 2017€)</i>				
Mean	3,300	1,600	800	800
St.D.	12,000	7,000	4,900	5,000
<i>e. Employment and participation rates (in %)</i>				
Employment	95.0	95.5	77.2	82.0
Participation	96.2	96.2	79.1	83.9

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Table 4.F.5 – Summary statistics: displaced workers' household income two years before separation (income tax returns)

Gender	Women		Men	
	Non-displaced	Displaced	Non-displaced	Displaced
<i>a. Disposable income (in 2017€)</i>				
Mean	46,400	46,900	45,800	46,700
St.D.	24,500	27,200	23,300	25,500
<i>b. Wages (in 2017€)</i>				
Mean	42,000	43,400	41,900	42,800
St.D.	25,600	29,200	24,000	26,800
<i>b. Unemployment benefits (in 2017€)</i>				
Mean	1,100	1,300	1,000	1,200
St.D.	3,100	3,500	2,900	3,100
<i>c. Other earnings (in 2017€)</i>				
Mean	3,700	3,400	3,200	3,400
St.D.	10,600	8,400	9,200	8,600
<i>d. Capital income (in 2017€)</i>				
Mean	4,600	4,500	4,400	4,500
St.D.	8,100	8,200	7,300	7,500
<i>e. Taxes (in 2017€)</i>				
Mean	-7,500	-8,100	-7,300	-7,800
St.D.	7,00	8,500	6,800	7,800
<i>f. Social benefits (in 2017€)</i>				
Mean	2,600	2,400	2,500	2,500
St.D.	3,700	3,500	3,800	3,800
<i>g. Position in the equivalent income distribution (in %)</i>				
First quintile	8.0	9.5	13.4	11.9
Second quintile	17.5	20.7	21.8	19.2
Third quintile	25.3	20.0	24.3	21.1
Fourth quintile	25.4	23.4	23.7	23.8
Fifth quintile	23.9	26.4	16.9	24.0

Source. Insee, DADS comprehensive files; Insee and Ministry of Finance (DGFIP), EDP sample.

Conclusion

Cette thèse développe quatre chapitres indépendants consacrés à l'effet de la vie familiale sur les décisions d'offre de travail en France. Au-delà de leurs objets connexes, ces chapitres sont également liés entre eux par leur recours commun à des données administratives individuelles et détaillées, produites et mises à disposition par le service statistique public. Ils partagent également l'idée, aujourd'hui commune en économie appliquée, de traiter comme une quasi-expérience des événements auxquels sont confrontées les familles – l'arrivée d'un enfant, l'expansion des services de crèche dans la commune, ou encore le licenciement d'un adulte – pour en inférer l'effet causal.

Le chapitre 1 montre que la diminution d'offre de travail des mères liée à l'arrivée d'enfants dans la famille, qui explique aujourd'hui la plus grande part des inégalités de genre sur le marché du travail, dépend largement de leurs salaires horaires potentiels. À l'inverse, les pères ajustent peu ou pas leur offre de travail, quels que soient les rendements du temps qu'ils passent en emploi salarié. Cette asymétrie suggère qu'alors que les familles prennent leurs décisions d'offre de travail en arbitrant entre les revenus générés par le temps passé par les mères sur le marché du travail, et les rendements non-monétaires du temps passé au sein du foyer et vraisemblablement consacré à l'éducation des enfants, elles ne prennent pas en compte un tel arbitrage en ce qui concerne les pères. Ainsi, si les normes de genre distordent l'allocation du temps et de l'effort au sein de la famille, ce n'est peut-être pas en prescrivant une quantité de temps et d'effort consacrée par les mères à l'éducation de leurs enfants : cela ne permet pas d'expliquer pourquoi les mères les mieux rémunérées choisissent d'y consacrer vraisemblablement moins de temps que les autres. À cet égard, mieux comprendre la très grande rigidité de l'offre de travail des hommes semblerait particulièrement profitable.

Le chapitre 2 poursuit cette étude des décisions d'offre de travail des parents de jeunes enfants en s'intéressant à l'effet de la disponibilité de places en crèche, dont les prix bas résultent d'importantes subventions publiques. Il montre que dans le cas français, ces services de garde de très jeunes enfants se substituent vraisemblablement moins à du temps parental, et en particulier maternel, qu'au recours à d'autres services de garde formelle fournis notamment par les assistantes maternelles. D'un point de vue normatif, ce résultat ne s'interprète pas nécessairement comme l'échec de cette politique publique. En effet, celle-ci doit également être évaluée du point de vue de la différence de qualité entre le service fourni par les crèches, et le mode de garde contrefactuel. En revanche, ces résultats attirent l'attention sur les mécanismes

d'attribution des places en crèche. En effet, si les familles auxquelles sont attribuées ces places appartiennent à l'ensemble de celles qui auraient dans le cas contraire recours à une assistante maternelle, alors il est possible qu'elles ne soient pas les familles qui en tireraient le plus grand bénéfice. Pour autant que l'on puisse en juger, des données quantitatives, ou même qualitatives, permettant de caractériser ces mécanismes d'attribution ne sont pas disponibles dans le contexte français. Leur mise à disposition semblerait pourtant particulièrement utile pour mieux comprendre les coûts et les bénéfices de cette politique publique.

Le chapitre 3 reprend en partie l'objet du chapitre 1, à savoir l'effet des enfants sur l'offre de travail de leur mère. Il vise en particulier à quantifier la contribution de cet effet à l'offre agrégée de travail infirmier, dans un contexte où celle-ci est réputée insuffisante. Si cet effet explique en partie – pour environ un tiers – la décroissance de l'offre de travail des infirmières hospitalières au long de leur cycle de vie, il ne permet en revanche pas de comprendre pourquoi une partie d'entre elles se retirent de l'emploi salarié après quelques années de carrière. Ce travail a vocation à être prolongé dans deux directions. D'une part, dans le cadre de l'étude de la profession infirmière et de sa participation aux systèmes de santé, la disponibilité prochaine de données longitudinales permettant de suivre à la fois la carrière dans l'emploi salarié, et dans l'emploi indépendant permettrait de quantifier la contribution des transitions vers l'exercice libéral au flux de sortie de l'emploi salarié. En effet, les implications de ces transitions quant à l'efficacité des systèmes de santé diffèrent profondément de celles de l'inactivité. D'autre part, l'approche mobilisée pour ce chapitre sera étendue à d'autres professions. En particulier, dans le contexte de l'emploi public en France, l'extension aux institutrices semblerait particulièrement instructive : elle permettrait de comparer deux professions très féminisées, aux niveaux de rémunération et de qualification proches, mais dont les conditions de travail diffèrent radicalement en particulier du point de vue de la flexibilité et du caractère prévisible des horaires. Cette comparaison contribuerait donc à une meilleure compréhension du rôle de la flexibilité horaire dans la conciliation entre vie familiale et vie professionnelle.

Le chapitre 4, enfin, s'intéresse aux conséquences du licenciement d'un salarié sur la vie familiale, et sur l'offre de travail des autres membres de sa famille. Ses résultats sont essentiellement négatifs : les décisions de mise en couple ou de séparation, la fécondité, ou encore l'offre de travail du conjoint ne semblent pas affectés par ce choc. À court terme, les salariés semblent relativement bien assurés contre la perte de revenus occasionnée par le licenciement, grâce à l'indemnisation du chômage. En re-

vanche, à long terme l'assurance chômage devient négligeable, l'essentiel des salariés retrouvant un emploi ; seulement, les revenus associés restent durablement inférieurs à ce qu'ils seraient en l'absence de licenciement, même après prise en compte de toutes les couches du système socio-fiscal. Ces résultats permettent de répondre à certaines questions soulevées par les chapitres précédents. D'une part, du point de vue méthodologique, l'absence d'effet de ces importantes pertes de revenu sur les décisions de fécondité suggère que le type de comparaison en différence-de-différences mobilisé par les chapitres 1 et 3 pour identifier l'effet causal des enfants n'est pas affecté par un biais d'endogénéité : dans le cas contraire, la perte d'emploi des femmes entraînerait vraisemblablement une augmentation de leur fécondité. D'autre part, ils suggèrent également que la spécialisation conjugale fondée sur l'avantage comparatif n'est pas toujours le cadre le plus satisfaisant pour comprendre les décisions jointes d'offre de travail au sein du couple : ce cadre prédirait en effet que le licenciement d'un des conjoints, en diminuant durablement la productivité du temps passé par lui sur le marché du travail, devrait occasionner un changement dans l'allocation du temps et de l'effort au sein du ménage, qui n'est pas observé dans les données. Ce travail a vocation à être prolongé dans deux directions différentes. D'une part, afin d'objectiver davantage le rôle assurantiel du système socio-fiscal français, la méthode développée dans ce chapitre sera étendue à d'autres chocs, par exemple de santé. D'autre part, un nouvel appariement entre les données administratives mobilisées dans ce travail et des données détaillées de consommation de soin doit permettre d'utiliser le licenciement comme un choc exogène sur les revenus des patients, ce qui autorisera notamment à étudier la prise en compte par les professionnels de santé des ressources financières de leurs patients.