Fertility and Labor Supply Responses to Child Allowances: The Introduction of Means-Tested Benefits in France

Elmallakh, Nelly

Demography (2023) 60(5):1493–1522 Published online: 21 September 2023
DOI 10.1215/00703370-10965926

Fertility and Labor Supply Responses to Child Allowances: The Introduction of Means-Tested Benefits in France

Nelly Elmallakh

ABSTRACT This article examines fertility and labor supply responses to a 2014 French policy reform that consisted of conditioning the amount of child allowances on household income. Employing regression discontinuity design and French administrative income data, I find that restricting family allowance eligibility criteria decreases fertility among the richest households. The results also highlight that receiving half the amount of the allowances or none leads to an increase in both male and female labor supply through an increase in overtime work. The implied change in earned income, due to an increase in weekly working hours, is found to be comparable to the euro value reduction in benefits. Auxiliary regression analyses show that the fertility decline reflects a decrease in the probability of having an additional child for parents rather than in the probability of becoming parents for households without children.

KEYWORDS Family policy • Child allowances • Fertility • Labor supply • France

Introduction

Declining fertility is an ongoing policy concern in many advanced economies. In the European Union (EU), for instance, between 1960 and 2016, average fertility dropped from 2.6 to 1.6 births per woman, a fertility rate that is now below replacement-level fertility (World Bank 2020). In response, several EU member states are currently implementing pronatal policies, using cash incentives for childbearing in an effort to curb declining fertility rates. Gaining a better understanding of the impact of recent family policy reforms on fertility and labor markets is thus of great importance.

France has long had relatively high fertility rates compared with other European countries and more broadly with other OECD members. In 2013, France spent a total of 2.9% of GDP on family benefits, well above the average investment level in OECD countries, which was 2.1% (OECD 2018). This high level of public spending reflects both higher than average child-related cash benefits and higher than average in-kind benefits for families with children within the OECD (OECD 2018, 2021). However, since 2010, the number of births in France has been declining. Between 2010 and 2018, the fertility rate in France fell from 2.03 to 1.87 births per woman. In-cash family benefits (as a percentage of GDP) also declined over the same period, from 1.6% to 1.4% (OECD 2018).

ELECTRONIC SUPPLEMENTARY MATERIAL The online version of this article (https://doi.org/10.1215/00703370-10965926) contains supplementary material.

This study focuses on a recent and important family policy reform in France and examines its impact on fertility and labor supply. The reform consisted of conditioning the amount of the basic allowances of early childhood benefits (Allocation de base des Prestations d'Accueil du Jeune Enfant, PAJE) on household income. The new law was first discussed in the National Assembly, the lower chamber of the French parliament, in March 2013, and officially entered into force on April 1, 2014. The reform did not have any explicit intended effect on fertility but rather aimed at changing the amounts of household-targeted public support by reducing the assistance paid to wealthier households.

The reform created a notched benefits schedule that phases out basic allowances from "full benefits" to "half benefits" and from "half benefits" to "zero benefits" by applying a two-year lagged definition of household taxable income. The 2014 reform defined two income thresholds for benefits: "half benefits" and "zero benefits." Households falling below the half benefits income threshold would be eligible for the total amount of basic allowances of early childhood benefits. Households with an income higher than the half benefits threshold but lower than the zero benefits threshold would be eligible for half the amount of basic allowances. Finally, households whose income exceeded the zero benefits threshold would no longer receive any basic allowances of early childhood benefits.[1]

Relying on this quasi-experimental design, I study the impact of the 2014 family policy reforms in France on fertility choices and labor supply using a regression discontinuity design. Several institutional details of the reform make this quasi-natural experimental setting particularly compelling. To start, benefit eligibility is based on a two-year lagged definition of taxable household income. Coupled with the short time between the initial discussion of the reforms and their final implementation, this implies that there is little room for direct manipulation of the running variable. Using data from the Statistiques sur les Ressources et les Conditions de Vie (SRCV; Statistics on Income and Living Conditions) for the years 2014 and 2015, I exploit the "sharp" discontinuity in the provision of the basic allowances of early childhood benefits to examine its impact on birth probability at the household level, as well as on the number of hours of work per week for both women and men after the reform.

This article is linked to a large body of academic literature on fertility responses to cash transfers and welfare reforms. Whether in the form of parental leave benefits, child allowances, or childcare subsidies, these studies have overall found that financial incentives have a positive effect on fertility; for instance, for Sweden, see Hoem (1993) and Björklund (2006); for the United States, see Rosenzweig (1999), Averett and Whittington (2001), and Joyce et al. (2004); for Canada, see Milligan (2005); for Austria, see Lalive and Zweimüller (2009); for Russia, see Malkova (2018); for western Europe, see Kalwij (2010); for the United Kingdom, see Brewer et al. (2012); for Germany, see Haan and Wrohlich (2011), Bauernschuster et al. (2016), and Raute (2019); for Israel, see Cohen et al. (2013); for Spain, see González (2013) and

[1] It is important to note that the income thresholds defined by the policy intervention take into account the number of children and the household structure. For example, a one-parent household with a working parent and one child is subject to a different income threshold than a two-parent household with one working partner and one child. See Table A3 in the online appendix for the different income thresholds defined by the 2014 reform.

González and Trommlerová (2023); and for France, see Piketty (2005) and Laroque and Salanié (2014).[2] Using data from the French Labor Force Surveys 1997, 1998, and 1999, Laroque and Salanié (2014) simulated the increase in births that would result from adding to the existing tax benefit system—an unconditional child benefit of 150 euros per month with a direct cost of 0.3% of GDP—and found that it might raise fertility by about 0.3 percentage points and reduce female labor supply by 0.5 percentage points.

On the other hand, a few studies have found no effect of benefits on fertility. For example, Kearney (2004) examined the impact of child benefit caps paid for an additional child in the United States and found no effect on fertility. Crump et al. (2011) found that U.S. tax benefits for children do not affect the level of fertility but may affect the timing of fertility, while Riphahn and Wiynck (2017) showed that an increase in the German child benefit did not change the fertility of low-income couples. More recently, researchers have been interested in uncovering the long-term effects of pro-fertility policies. Parent and Wang (2007), Kim (2014), and Adda et al. (2017) suggested that welfare policies have little long-term effect on fertility and generate only timing effects, as they induce women to have children earlier.

My study also relates to the literature on the labor market impact of financial incentives for welfare recipients.[3] Parental leave coverage was found to be associated with higher women's employment (Ruhm 1998) and with large increases in mothers' time away from work after birth (Baker and Milligan 2008). It also appeared to have a strong impact on mothers' return to work after childbirth (Berger and Waldfogel 2004; Dustmann and Schönberg 2012) but to have no effect on women's wages (Albrecht et al. 1999; Baum 2003).[4]

The present study contributes to the literature in at least three ways. First, while most of the above-cited literature has focused on fertility responses to increases in monetary incentives provided in the form of various welfare reforms, this paper contributes to the rather small literature on the impact of the elimination of welfare reforms or caps in the provision of monetary incentives on fertility.

Second, the natural experiment that allows for causal identification is unique and advantageous, as eligibility in a given year was legislatively determined by household income two years before. For instance, in 2014, when the means-tested PAJE benefits were introduced, eligibility was determined by taxable household income in 2012. In this setting, the placement of families on either side of the threshold is determined before the threshold is actually introduced—therefore families cannot influence their eligibility by altering their income. The methodology used here also allows the exploitation of differences in entitlements across households, in contrast to the

[2] See Gans and Leigh (2009) on the effects of the announcement of a "baby bonus" in Australia and Brunner and Kuhn (2014) on the effects of the announcement of the abolition of a baby bonus in Austria. In both studies, the authors showed that fertility responded to the announcement of these two policy changes before their actual implementation or abolition.

[3] This study is also linked to a broader literature on the effect of welfare on living arrangements. For example, see Card and Lemieux (1997) on the United States and Canada.

[4] It is important to note that in Baum (2003), the analysis focused on the Family and Medical Leave Act (FMLA) in the United States in 1993 and that the policy guaranteed 12 weeks of unpaid leave for eligible mothers. The author argued that the short and unpaid nature of the maternity leave could explain the results.

existing literature, which for the most part focused on the generalized introduction or extension of welfare policies.

Third, one of the main novelties of this study is the use of the SRCV data. The SRCV dataset has the unique feature of providing administrative information on household income, thus providing an additional contribution with respect to the existing literature, which has mostly employed survey data. Finally, given the importance of this historical shift in welfare state design as France broke for the very first time with the concept of "benefit universality," it is critical to understand the effect of these recent reforms on fertility responses and labor supply.

My results suggest that, among middle-income couples, cutting the early childhood benefits by half does not have a significant impact on fertility. On the other hand, among the richest households, the complete removal of the benefits leads to a decline in fertility.[5] I also explore whether the effect of the policy kicked in when the policy was first discussed, in March 2013, or at the time the policy went into effect in April 2014. The former definition is employed in the benchmark model—as it reflects the outer bound on when any birth effect could possibly occur—and leads to an estimated benefit elasticity of 0.38, while the latter definition results in a smaller implied elasticity of 0.13. This suggests that the largest decline in fertility occurred following the earliest discussions of the policy reform in the lower chamber of parliament, a process that received substantial media coverage, as I will later highlight in further detail.

The findings also suggest that the reform had no impact on the decision of becoming parents for households without children prior to the reform, but did affect the decision of having an additional child for parents. Indeed, I find that the decline in fertility, following the announcement and the implementation of the reform, affected only the probability of having a second child or children of higher birth parity.

How do these effects compare to those of other studies? My estimates are within the range of previously estimated elasticities of fertility to benefits or to tax exemption in the literature. For instance, Gauthier and Hatzius (1997) estimated a long-run, cross-country benefit elasticity of 0.16. For the United States, Whittington et al. (1990) estimated the elasticity of fertility to personal tax exemption for dependents to be between 0.13 and 0.25. Whittington (1992) subsequently found larger elasticities, ranging from 0.23 to 1.19. In the case of Quebec, Milligan (2005) estimated a benefit elasticity of 0.11, whereas for Israel, Cohen et al. (2013) estimated an overall benefit elasticity of 0.19 (higher elasticities were observed among certain religious groups, reaching as high as 0.33). As for Europe, Laroque and Salanié (2014) found an elasticity of fertility to tax benefits of 0.20 in France, while Brewer et al. (2012) estimated an elasticity of 0.28 in the United Kingdom.

Regarding the impact on labor supply, I find that receiving half the amount of basic allowances of early childhood benefits increases the number of hours of work per week by two to four hours, for both women and men, compared with those who receive full benefits. Likewise, being ineligible for any basic allowances also increases the number of hours of work per week for women and men, relative to

[5] This finding is in line with Joyce et al. (2004), who found that, in the United States in the 1990s, birth rates fell in states that implemented provisions to reduce cash assistance to welfare recipients who have additional children.

individuals who are eligible for either full or half the allowances.[6] A back-of-the-envelope calculation shows that the implied change in earned income, due to an increase in weekly working hours, compares to the euro value reduction in benefits. Assuming average net hourly wages in the year 2014, derived from the Institut National de la Statistique et des Études Économiques (INSEE; French National Institute of Statistics and Economic Studies), a two-hour increase in the number of hours of work per week generates approximately 95 euros per month for women and 113 euros per men, covering a 50% reduction in monthly allowances (i.e., a loss of 92.31 euros per month). On the other hand, a four-hour increase in the number of weekly working hours increases women's monthly earnings by 190 euros and men's by 226 euros. This income increase fully covers a 100% cut in the monthly basic allowances of early childhood benefits, which amounts to a loss of 184.62 euros per month.[7]

Theoretical Framework and Background

The Economic Incentives of the Benefits

Becker (1960) argued that parents should be viewed as rational economic actors and children as durable consumption and production goods. Becker's model stresses the importance of the cost of children in explaining fertility differentials. Accordingly, the demand for children responds to changes in the price of the marginal child. The framework implies that reductions in child subsidies would lead to an increase in the price—net of subsidy—of the marginal child, and hence to a decline in the demand for children, as a result of the optimizing behavior of parents and would-be parents.[8]

Becker-type models of fertility are controversial among demographers, as highlighted by Olsen (1994) and Milligan (2005). Indeed, standard demographic analysis tends to emphasize the role of social norms, biological processes, and reproductive technology in explaining fertility and would therefore predict no changes in fertility in reaction to exogenous price changes.

The PAJE reform affected child allowance receipts for richer households but not for the poorest households (those whose income falls below the half benefits threshold). As for households whose income exceeds the half benefits threshold but is below the zero benefits threshold, they experienced a reduction in PAJE receipts by half, while households whose income exceeds the zero benefits threshold experienced a total elimination of these allowances. I test whether, in line with the predictions of Becker's model, the decline in child subsidies following the 2014 reforms was associated with reduced births. The finding that fertility responds to price changes does in

[6] This is in line with Moffitt (2002), who provided a survey of the existing literature on the labor market incentive effects of transfer programs in the United States and concluded that the elimination of welfare programs would lead to an increase in the hours of work by 10% to 50%.

[7] According to the INSEE (Arnault 2018), the average gross hourly wage for men in 2014 is 18.8 euros and for women is 15.8 euros. The difference between gross and net salaries in France is generally between 20% and 30%. In the estimations above, I employ an average rate of 25%.

[8] Hotz et al. (1997) provided a detailed discussion of fertility models, surveying the theoretical and empirical literature.

fact provide evidence in favor of Becker's model and, hence, supports the theory that prices matter in explaining fertility choices.

The impact of the reform on parental labor supply, on the other hand, can be discussed within a simple labor supply model. Schøne (2004) posited that parents adjust their labor supply as they maximize the value of consumption and leisure subject to a budget constraint. If a household is eligible for child allowances, the budget constraint will be positively affected by the allowances. Given the fungibility of income sources, whether income comes in the form of child allowances or in the form of any other income source, what matters in determining parental labor supply is the total sum of all income components.

Within a simple leisure–consumption framework, an increase in household income translates into a positive shift in the budget constraint, leading to an increase in the leisure–consumption combination. Standard labor supply theory predicts that the child benefits introduce both substitution and income effects that act in the same direction toward reduced labor supply if leisure is a normal good. For working parents, the reform will generate both income and substitution effects. For parents who are not working, the PAJE benefits introduce only an income effect. In the case of PAJE benefits, the eligibility criteria stipulate that at least one of the two parents is working. The decrease in parental labor supply would therefore depend on the total amount of child benefits, as well as on the parents' preferences for leisure and consumption.

The 2014 Family Policy Reform

The French family policy system follows a long-established pronatal and family-centered tradition. Compared with other OECD countries, France has one of the most generous schemes for family benefits and spends relatively more in terms of public investment in families with children (OECD 2018).

In 2014, the French government undertook a series of social and family policy reforms that aimed at changing the amounts of household-targeted public support. The reforms proposed several measures that led to a reduction in the assistance paid to wealthier families while increasing transfers to the most vulnerable households. This article focuses in particular on the 2014 reform to the basic allowances of the early childhood benefits (Allocation de base des Prestations d'Accueil du Jeune Enfant). The PAJE includes a package of benefits to compensate for the cost of raising children: a birth or adoption premium (Prime de naissance ou d'adoption), basic allowances (Allocation de base), and allowances to support the reconciliation between personal and professional life and to compensate for childcare costs.

The basic allowances aim to help households cover the costs of child education and maintenance. They are intended for parents of a child younger than three years old and are paid on a monthly basis for three years, from the first day of the month following the birth until the month the child turns three. In the case of multiple births, households may accumulate several basic allowances for their children. Several conditions determine eligibility for the basic allowances: (1) the child should be younger than three years old; (2) in the case of two-parent households, either one or both parents should have been working and generating income two years prior to benefit receipt, and in the case of one-parent households the single parent should have been working and generating income two years prior to benefit receipt;[9] and (3) household income (as reported two years before) should have been lower than the income threshold defined by the 2014 policy reform.

The reform stipulated that for all children born or adopted as of April 1, 2014, working parents would be eligible for differential benefit rates depending on their disposable household income two years before (a pre-reform landscape is provided in online appendix A). Even though approximately 13 months elapsed between the beginning of the policy-making process and the final implementation of the reform, substantial media coverage over this period brought the basic allowances of PAJE benefits to the attention of the general public. Several news outlets, including *Le Monde* and *Le Figaro*, covered the policy reform process from the earliest discussions in the lower chamber of parliament. Articles published in April 2013 outlined, in general terms, the potential impact of the reforms on the allocation of PAJE benefits. In June 2013, another flurry of news articles explained in detail the full scope of the reform ("Allocations familiales" 2013; Guichard 2013; Laurent and Laurent 2013; "Le gouvernement choisit d'abaisser" 2013; Vignaud 2013).

The Google trends analysis depicted in Figure 1 confirms the surge in web search popularity in France for the term "paje," which refers to the Prestations d'Accueil du Jeune Enfant. In this figure, weekly web searches are averaged by month. Google trends search scores are plotted on the *y*-axis, representing search interest relative to the highest point on the chart for the given region and time. A value of 100 is the peak popularity for the term, and a value of 50 means that the term is half as popular.[10] Search popularity for PAJE benefits in France gained immediate momentum following the discussion of the law in the National Assembly in March 2013. While search popularity also increased around the time of the official law enactment in April 2014, analysis shows that searches peaked earlier when the law was first discussed in the National Assembly—likely due to the extensive media coverage the reform proposal received.

Despite media coverage of the policy-making process, the early news articles were uncertain and speculative about a potential reform underway, highlighting that the PAJE benefits might be means-tested. Articles published later on, starting from June 2013, outlined that the benefits would be conditional on income, but that the applicable income thresholds were not yet official. Indeed, thresholds were defined by the 2014 Financing Act, published on December 30, 2013, after several rounds of discussions and amendments between September and December 2013.[11] Section D1 in the

[9] According to the Caisse d'Allocations Familiales (CAF), the income considered is generated from professional activity. The only exception is in the case of occupational diseases or work accidents, for which the income considered is the daily allowances for accidents at work or occupational diseases. A parent's income should be at least equal to 5,252 euros two years before the survey.

[10] It is important to note that the Google trends search score data are weekly, while Figure 1 relies on monthly averages. Therefore, the value of 100, which corresponds to the peak of popularity of the term over the entire period, does not appear in the figure. In fact, the term "paje" was most popular in March 2014 before the actual implementation of the reforms in April 2014. Both weekly and monthly data confirm that the search popularity of PAJE benefits gained momentum long before the actual enforcement of the reforms in April 2014.

[11] The official website of the French Senate provides more information on the 2014 Financing Act and the various amendments brought about by the National Assembly and the Senate between September and December 2013 (https://www.senat.fr/dossier-legislatif/pjlf2014.html).

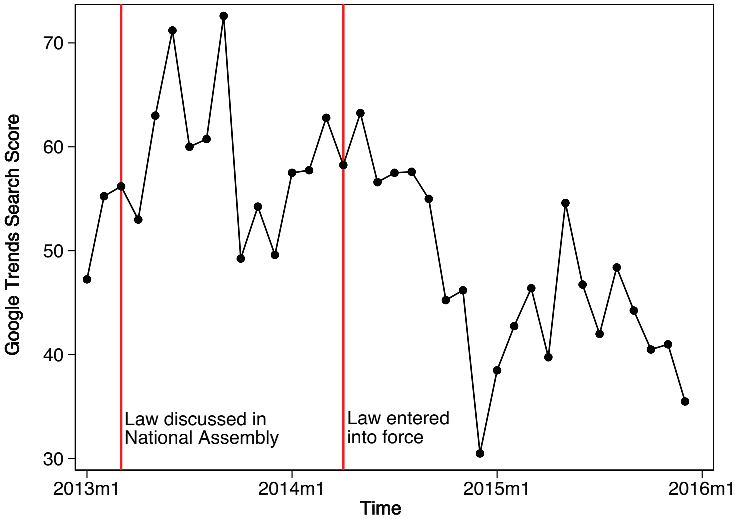

Fig. 1 Google trends analysis of the PAJE reforms, showing web searches in France for the term "paje," which refers to the Prestations d'Accueil du Jeune Enfant. The vertical red lines correspond to March 2013, when the law was first discussed in the National Assembly, and to April 2014, when the law came into force. Values on the *y*-axis represent search interest relative to the highest point on the chart for the given region and time. The underlying raw data are weekly, with a value of 100 for the peak popularity for the term, while a value of 50 means that the term is half as popular.

online appendix shows that there is no evidence of bunching in household income in 2013 on either side of the discontinuity thresholds (all sections, tables, and figures designated with an "A" are available in the online appendix).

The 2014 policy reform defined income thresholds that depend on the number of children and household structure (see Table A3). The thresholds applicable to a one-parent household with income are identical to those for a two-parent household with two incomes.

Prior to the 2014 policy reform, working parents were eligible for a universal amount of basic allowances, unconditional on household income. Following the reform, depending on household income, the poorest households were still eligible for the total amount of basic allowances (184.62 euros per month), middle-income households were eligible for half the amount (92.31 euros per month), and the wealthier households, whose income exceeded the designated threshold, were no longer eligible for any basic allowances.[12] It is important to note that the total amount of the basic allowances is the same as the previous universal benefit that was applicable to all working parents.

[12] The loss of child allowances as a percentage of household income can be computed using the average household income for households that receive full benefits, half benefits, and zero benefits. A back-of-the-envelope calculation shows that these losses may not be negligible. Indeed, with only one child younger than three, the loss of allowances would be equivalent to 2.4% and to 3% of yearly income for households with half benefits and zero benefits, respectively. With two children younger than three, this loss doubles to 4.8% and 6%, respectively.

Data

This study relies on data from the Statistiques sur les Ressources et les Conditions de Vie. The SRCV dataset corresponds to the French part of the European Union–Statistics on Income and Living Conditions (EU-SILC). The Institut National de la Statistique et des Études Économiques implements the SRCV surveys, which have been conducted in France on a yearly basis since 2004 and provide representative data on income and living conditions in both cross-sectional and longitudinal dimensions. Each year, approximately 12,000 households are surveyed, including a refresher sample of nearly 3,000 households. As in a typical household survey, the SRCV dataset includes household as well as individual questionnaires and covers topics such as family composition, housing and living conditions, household income, taxes, social security costs incurred by households, and social benefits. The individual-level questionnaire provides information on educational background, economic activity, occupation, income, and labor supply, among other variables.

A unique feature of the SRCV dataset is that it collects information on household income and social benefits from various administrative sources instead of collecting them from survey respondents. The (taxable) income of surveyed households is derived from their tax returns. The social benefits that households eventually receive are obtained from the governing body on which the households depend: the National Family Allowance Fund (Caisse Nationale des Allocations Familiales), the Agricultural Social Mutual Fund (Mutualité Sociale Agricole), or the National Pension Insurance Fund (Caisse Nationale d'Assurance Vieillesse).

The inclusion of this administrative data constitutes one of the greatest advantages of the SRCV dataset relative to typical individual and household surveys, in which respondents could misreport or manipulate their income information.[13] The population of interest consists of households with children and, at the time of benefit receipt, the child must be younger than three years old. In relying on data from the 2014 and 2015 survey rounds, I focus on N-children households ($N > 0$) in each survey round, compute their distance to thresholds for N-children families, and investigate whether a household decides to increase its size from $N - 1$ children (prior to the reform) to N children (at the time of the survey). This allows me to investigate the impact of potential future allowances on both the probability of becoming parents for those without children at baseline and the probability of having one additional child for those who are already parents. In both cases, if a household decides to have a child in the window under consideration, the child is by definition aged less than three.

Given the policy reform, the treatment status of N-children households is determined by their household income two years preceding the survey year. In any given year in the SRCV survey, the survey is matched with administrative income data of the year before. The 2014 survey therefore collects household income in the year 2013, while the 2015 survey collects household income in the year 2014. To obtain income data two

[13] While the SRCV surveys are advantageous in many ways, the data collected on family and child allowances are not disaggregated. Indeed, the surveys provide information on whether households receive family or child allowances without distinction between the various types of benefits that a household could receive. Therefore, it is not possible to check the discontinuity of basic allowances of early childhood benefits in the neighborhood of the discontinuity thresholds.

years prior to the survey, for the year 2014, I restrict my analysis to households that are also observed in 2013 in order to be able to collect household income in the year 2012. In a similar way, for the year 2015, the analysis is restricted to households that are also observed in 2014, which allows the collection of information on household income in 2013, two years preceding the survey. The estimation sample therefore consists of pooled cross-sectional household data for the years 2014 and 2015, with income information for the years 2012 and 2013, respectively.[14] The SRCV data also include information on the number of children in a given household, as well as the year and month of birth of the children born in both the survey year and the preceding one. Therefore, the data allow identifying the exact number of children a household has at the time of benefit receipt, the two-year lagged definition of taxable income, whether a household is two-parent or one-parent at baseline—in line with the two-year lagged income definition—and the family's labor force participation. These variables allow one to identify the various household groups (presented in Table A3), making it possible to accurately determine the precise income threshold that is applicable to each family.

In my estimation sample, 47% of households fall below the half benefits threshold, 16% are in between the half benefits and the zero benefits thresholds, and 37% fall above the zero benefits threshold. Table 1 presents descriptive statistics at the half benefits and zero benefits thresholds for women and men. These statistics are restricted to individuals who fall within a range of 5,000 euros below and above the income thresholds defined by the 2014 family policy reform. Individual characteristics include the individual's age, nationality, and education.[15] This table confirms random sorting around the discontinuity thresholds and shows that individual covariates are balanced in the neighborhood of the thresholds.[16] Appendix section D2 also provides descriptive statistics at these thresholds at the household level, as well as covariate balance tests to check whether the treatment status can be considered as good as randomly assigned. Appendix section D2 also shows that there is not any statistically significant difference between individuals below and above the discontinuity thresholds along these multiple covariates.

[14] Even though I rely on the longitudinal dimension of the data to retrieve income information, the data analysis remains cross-sectional in nature since I do not focus exclusively on households that were observed simultaneously in 2014 and 2015. Such a focus would significantly reduce my sample size, which would require households to be observed over three consecutive survey years (in 2013, 2014, and 2015). Instead, my analysis also includes households interviewed in 2014 but not necessarily in 2015 (as long as they were interviewed in 2013 to retrieve their income information).

[15] Unfortunately, the SRCV dataset does not provide information on parental leave take-up or parental leave duration, which would allow one to examine such behaviors in the neighborhood of the discontinuity thresholds. Nonetheless, according to the most recent data on family benefits from the French Social Security—which provide detailed information on the beneficiaries of parental leave in France (*la Prestation partagée d'éducation de l'enfant, PreParE*)—the highest proportion of the PreParE beneficiaries belongs to the middle-income deciles. Indeed, 29% of the beneficiaries belong to the fourth and fifth income deciles, while lower proportions of individuals belong to the lower and upper ends of the income distribution benefit from the PreParE. This would suggest that parental leave take-up is likely to be low at both the lower and the upper ends of the distribution. These statistics and the evidence provided in appendix section B3 suggest that parental leave take-up is not driving my results.

[16] According to the United Nations Economic Commission for Europe (UNECE), the mean age of women at birth of the first child in France was 28.3 in 2014 and 28.4 in 2015. The mean women's age in the sample around the discontinuity thresholds is roughly consistent with these figures, in particular for women who belong to the wealthiest households (those above the zero benefits threshold).

Table 1 Individual-level descriptive statistics (means) at the discontinuity thresholds

	A: Women				B: Men			
	Half Benefits		Zero Benefits		Half Benefits		Zero Benefits	
	Below (1)	Above (2)	Below (3)	Above (4)	Below (5)	Above (6)	Below (7)	Above (8)
Age	26.890	26.480	26.780	28.430	27.350	27.700	28.230	28.400
	(17.000)	(16.930)	(17.000)	(17.320)	(17.430)	(17.780)	(17.990)	(17.540)
Number of observations	776	726	667	560	732	694	635	572
Nationality								
French	.970	.982	.979	.979	.970	.973	.981	.982
	(.171)	(.131)	(.144)	(.144)	(.170)	(.163)	(.138)	(.134)
EU citizen	.010	.011	.014	.013	.013	.016	.015	.010
	(.100)	(.104)	(.118)	(.115)	(.112)	(.125)	(.120)	(.101)
Algerian, Moroccan, or Tunisian	.008	.000	.000	.003	.002	.000	.000	.002
	(.089)	(.000)	(.000)	(.052)	(.046)	(.000)	(.000)	(.051)
African except Maghreb	.006	.000	.002	.000	.006	.002	.002	.000
	(.078)	(.000)	(.048)	(.000)	(.080)	(.048)	(.049)	(.000)
Other nationality	.006	.007	.005	.005	.008	.009	.002	.005
	(.076)	(.081)	(.068)	(.073)	(.092)	(.095)	(.049)	(.072)
Number of observations	405	457	428	377	471	443	413	386
Education								
First cycle education or less	.198	.145	.149	.144	.251	.193	.157	.183
	(.399)	(.352)	(.357)	(.352)	(.434)	(.395)	(.364)	(.387)
Second cycle education	.553	.492	.450	.442	.560	.511	.523	.483
	(.498)	(.501)	(.498)	(.497)	(.497)	(.501)	(.500)	(.500)
Postsecondary education	.148	.205	.221	.240	.128	.142	.172	.173
	(.356)	(.404)	(.415)	(.428)	(.335)	(.349)	(.378)	(.379)
University or above	.101	.158	.181	.173	.060	.155	.148	.161
	(.302)	(.366)	(.385)	(.379)	(.238)	(.362)	(.356)	(.368)
Number of observations	405	366	349	312	398	374	331	323

Notes: Means are shown, with standard deviations in parentheses. The analysis is restricted to households within ±5,000 euros of yearly household income from the income thresholds defined by the 2014 policy reform. "Below" corresponds to 5,000 euros below the threshold and "above" corresponds to 5,000 euros above the threshold. First cycle education or less includes pre-primary, primary, and first cycle of basic education; second cycle education includes second cycle of basic education, of general secondary education, or of vocational secondary education; postsecondary education includes postsecondary nonuniversity education and short cycle university education; and university education or above includes bachelor's, master's, and Ph.D. levels.

Empirical Strategy: A Regression Discontinuity Design

The empirical analysis relies on a regression discontinuity design (RDD) to examine the impact of the 2014 family policy reforms on fertility and labor supply for women and men. Within the framework of a "sharp" regression discontinuity, the treatment status is a deterministic function of a continuous variable: the forcing variable or the assignment variable. Individuals therefore receive or do not receive the treatment according to the underlying value of the forcing variable, as illustrated by[17]

$$T_i = T(X_i) = 1[x_i \geq \dot{x}]. \tag{1}$$

$1[\cdot]$ is an indicator function, x_i is the forcing variable, and $\dot{x}$ is the discontinuity threshold, which is the value taken by the forcing variable separating the units into two mutually exclusive groups, that is, those who receive the treatment versus those who do not receive it.

A unique feature of the policy reform in France is the differential benefit rate by income. If the household income is below a certain threshold $\overline{R}$, households are eligible for full benefits. Meanwhile, households whose income is between $\overline{R}$ and $\overline{\overline{R}}$ are eligible for half benefits. Finally, households whose income exceeds $\overline{\overline{R}}$ are eligible for zero benefits.[18] The analysis therefore involves two discontinuity thresholds: the first threshold ($\overline{R}$) will be denoted as "half benefits" and the second threshold ($\overline{\overline{R}}$) as "zero benefits." In Eq. (2), the mutually exclusive groups are those who receive full benefits versus those who receive half benefits:

$$T_{\overline{R}_i} = T(R_i) = 1[R_i \geq \overline{R}] \times 1[R_i < \overline{\overline{R}}]. \tag{2}$$

In Eq. (3), the analysis compares households/individuals who receive zero benefits with those who receive either full or half benefits:

$$T_{\overline{\overline{R}}_i} = T(R_i) = 1[R_i \geq \overline{\overline{R}}]. \tag{3}$$

As presented in Table A3, the income thresholds defined by the policy are a function of both the number of children and household structure. One of the main advantages of the RDD framework is the local quasi-randomization around the specific thresholds employed in the analysis. The RDD approach therefore consists in comparing the outcomes of households/individuals who are "just below" and "just above" the threshold. The intuition behind this approach is that households/individuals whose incomes are close to the discontinuity threshold are very comparable along observable and unobservable characteristics, except for the treatment. In other words, in the neighborhood of the discontinuity threshold, households are very similar and only some are subject to treatment. Therefore, the households/individuals "slightly below" the threshold provide the counterfactual for those "slightly above" the threshold since the treatment (receiving half or zero benefits) is effectively randomized in the neighborhood of the discontinuity thresholds. By extension, this quasi-randomization

[17] See Lemieux and Milligan (2008) and Pettersson-Lidbom (2008) for a similar methodology.

[18] It is important to note that $\overline{R} < \overline{\overline{R}}$.

around the discontinuity thresholds implies that there would be no reason to believe that the results estimated in the neighborhood of the discontinuity thresholds capture the effects of other economic factors—which are, unlike the reform in question, not randomized around the discontinuity thresholds but are rather common shocks.[19] The inclusion of year fixed effects in Eq. (4) captures common economic shocks that would affect households in the years considered.

The RDD can be implemented in several ways. The simplest is a nonparametric technique that compares outcomes in a small neighborhood below and above the discontinuity threshold. However, this approach could lead to imprecise estimates of the treatment effect, with the usual trade-offs in the choice of the bandwidth (Lemieux and Milligan 2008). When using a small bandwidth, the treatment effect is equal to the difference in the average outcomes of units that are "just below" and "just above" the threshold. This could lead to imprecise measures of the treatment effect. By contrast, using large bandwidths could lead to biased estimates of the treatment effect if units that are further away from the discontinuity threshold are systematically different from those around the discontinuity point. Hence, unless large sample sizes are available in the neighborhood of the discontinuity threshold, the nonparametric bandwidth RDD estimation is likely to be subject to a large degree of sampling variability.

Using a control function is an equivalent yet more efficient technique of estimating the treatment effect using RDD (Lemieux and Milligan 2008; Pettersson-Lidbom 2008). This approach balances the trade-off between precision and bias by using all the available data around the discontinuity threshold and regressing the outcome of interest Y_{it} on the treatment indicator T_{it}, the control function ($R_{i,t-2}$, which is a low-order polynomial of the treatment-determining covariate R_i), and the interaction term between the treatment indicator and the control function. The RDD specification is presented in the equation[20]

$$Y_{it} = \pi T_{it} + \delta(R_{i,t-2}) + \mu T_{it} \times \delta(R_{i,t-2}) + \lambda_t + \varepsilon_{\text{it}}. \tag{4}$$

The regression discontinuity is "sharp" in outcomes, since the outcome variables are measured at the time of the survey. The dependent variables Y_{it} correspond to the birth probability at the household level (for births that occurred nine months after the reform was first discussed in the National Assembly or implemented), as well as to the number of hours of work per week at the time of the survey, for women or men. The control function $\delta(R_{i,t-2})$ is a first-order polynomial of the forcing variable, which is the household income two years prior to the survey year and is equal to the difference between the total disposable household income net of contributions and

[19] This feature of RDD implies that it is very unlikely that households/individuals whose income is close to the discontinuity threshold would have differential responses to common economic shocks that would be driving the results. Indeed, we would expect households/individuals located below and above the discontinuity threshold to be affected by common shocks in a similar way. The placebo regressions included in appendix section E6 further corroborate this evidence. Indeed, if other economic factors, correlated with different income brackets, were driving the results, they would have been picked up in the placebo regressions reported in Table A19, which relies on placebo thresholds.

[20] Equation (4) denotes the individual-level estimation. However, household-level regressions are also estimated when examining the impact of the 2014 family policy reforms on fertility choices.

the discontinuity threshold defined by the policy. The regression specification also includes a year fixed effect λ_t, while ε_{it} denotes the error term.

Two sets of regressions are estimated separately. The first one compares the outcomes of individuals/households who receive half benefits with those who receive full benefits, as illustrated in Eq. (2), by restricting the analysis to those whose income is strictly below $\overline{\overline{R}}$. The second one compares the outcomes of individuals/households who do not receive any basic allowances of the PAJE benefits with those who either receive half or full benefits. The main coefficient of interest, π (reported in the regression tables), is the RDD treatment effect, and it measures the difference in average outcomes between treated and untreated units. More specifically, it measures the difference in outcomes between those who receive half benefits versus those who receive full benefits; and the difference in outcomes between those who receive zero benefits versus those who receive half or full benefits. The control function technique thus yields to unbiased estimates of the treatment effect, since $R_{i,t-2}$ is the only systematic determinant of the treatment status. The inclusion of the control function $\delta(R_{i,t-2})$ thus captures any correlation, which may otherwise occur, between $R_{i,t-2}$ and the error term ε_{it}.

The control function is my benchmark specification. Nonetheless, I also report results using nonparametric bandwidth estimations, since the estimates from the control function and the nonparametric bandwidth estimations should be the same if the control function is correctly specified. Moreover, in the online appendix, I present results including a vector of predetermined individual and household control variables. The provision of family allowances for children should not be systematically correlated with any observed or unobserved variables once the forcing variable is controlled for. Hence, adding additional control variables should not affect the estimates from the control function approach and that provides a test of whether the treatment status can be considered as good as randomly assigned (Lemieux and Milligan 2008).

Empirical Findings

A Graphical Presentation

Panel a of Figure 2 presents a regression discontinuity (RD) plot of the relationship between the birth probability and the distance from the zero benefits threshold. The distance to the discontinuity threshold is measured by the difference between a two-year lagged definition of household income and the income threshold defined by the 2014 policy reform. The vertical red line corresponds to the income threshold beyond which households are not eligible to receive the basic allowances of PAJE benefits. Positive values on the x-axis refer to households whose income is above the discontinuity threshold and thus receive no benefits, while negative values refer to households whose income is below the discontinuity threshold and are eligible for full or half benefits.

I define the probability of birth at the household level as a dummy variable indicator for births that occurred in the years 2014 and 2015, nine months after the law was

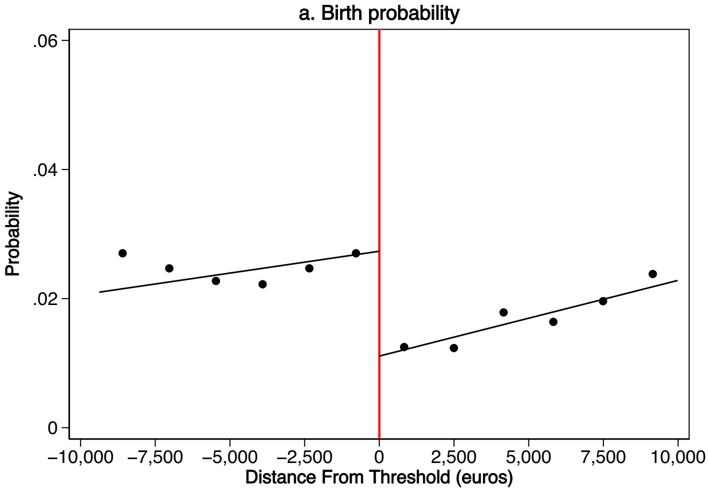

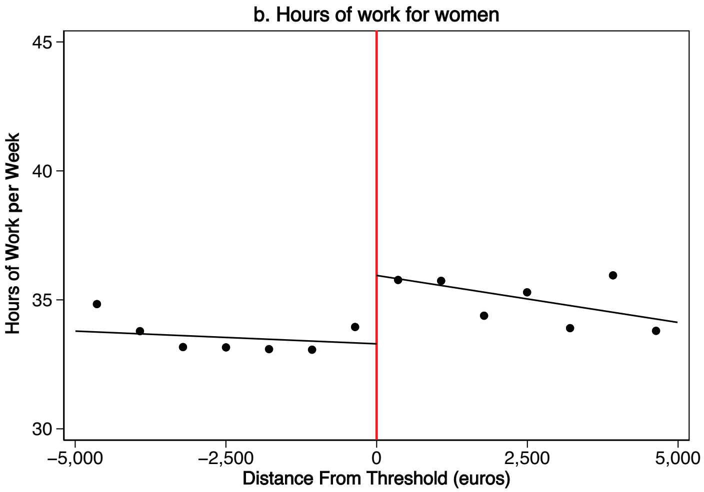

Fig. 2 Discontinuity at the "zero benefits" threshold. On the *y*-axes, the birth probability at the household level is plotted in panel a using household-level data, while women's and men's hours of work per week are plotted in panels b and c using individual-level data. The probability of birth at the household level is a dummy variable indicator for births reported nine months after the law was first discussed in the National Assembly in March 2013. On the *x*-axes, the distance to the "zero benefits" threshold is shown in relation to the vertical red line. It is computed as the difference between the income received by the household two years preceding the survey minus the income threshold defined by the 2014 policy reform. The red lines correspond to the income threshold beyond which households are not eligible to receive PAJE benefits. Positive values on the *x*-axes refer to households above the income threshold eligible for zero benefits; negative values refer to households below the income threshold eligible for half benefits.

(continued)

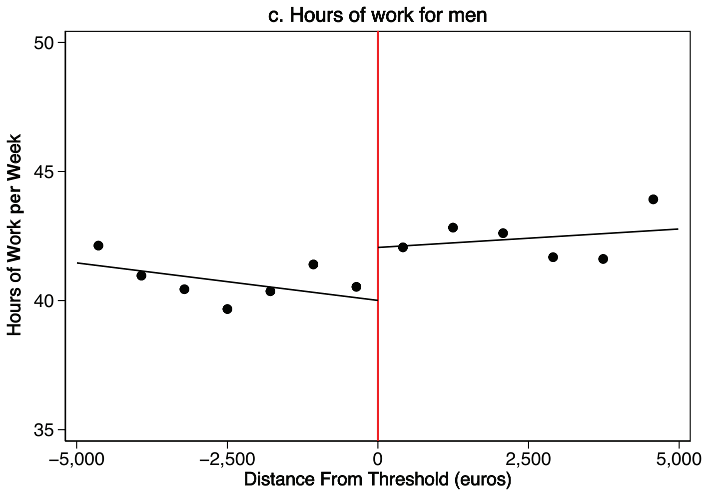

Fig. 2 (continued)

first discussed.[21] Figure 2 shows that households whose income exceeded the discontinuity threshold of zero benefits had a consistently lower birth probability compared with households that were eligible for half the amount of basic allowances of early childhood benefits.[22]

Regarding labor supply, panels b and c of Figure 2 show the discontinuity between the number of hours of work per week and the distance from the zero benefits threshold. These RD plots illustrate a clear discontinuity in the number of hours of work between the two groups, for both women and men. Receiving zero benefits was associated with a greater number of weekly working hours for women and men relative to those who were eligible for half the amount of basic allowances.

Control Function Specification

Turning to the results from the control function, Table 2 presents the RDD estimation results on the impact of the 2014 reforms on fertility (panel A) and labor supply (panel B). Equation (4) is estimated using a first-order polynomial specification of the control function on household data in panel A and individual data in panel B. In columns 1 and 2, the dependent variable is a dummy variable indicator for the probability of birth at the household level, for births that occurred nine months after the law was first discussed in the National Assembly. In columns 3 and 4, the dependent

[21] In Table 2, I also explore the impact of the 2014 reforms on births that occurred nine months after the law entered into force in April 2014, instead of nine months after the law was first discussed in the National Assembly.

[22] Restricting the analysis to households within narrower income windows results in noisier graphs, as the birth probabilities within many income bins are equal to zero. However, I show regression results in Table 3 using nonparametric bandwidth estimations with narrower household income windows of ±2,500, ±5,000, ±7,500, and ±10,000 euros.

Table 2 Control function specification: Impact of the PAJE reform on fertility and labor supply

	A: Impact on Fertility			
	Announcement Effect		Implementation Effect	
	Birth Probability (1)	Birth Probability (2)	Birth Probability (3)	Birth Probability (4)
Half Benefits	–.004		.006	
	(.018)		(.009)	
Zero Benefits		–.020*		–.007†
		(.008)		(.004)
Number of Observations	1,890	2,786	1,890	2,786
R^2	.027	.025	.019	.015
Year Fixed Effects	Yes	Yes	Yes	Yes

	B: Impact on Labor Supply			
	Announcement Effect		Implementation Effect	
	Hours of Work/Week for Women (1)	Hours of Work/Week for Women (2)	Hours of Work/Week for Men (3)	Hours of Work/Week for Men (4)
Half Benefits	3.575**		2.306*	
	(1.132)		(1.004)	
Zero Benefits		1.834*		4.088**
		(0.848)		(0.688)
Number of Observations	1,405	2,194	1,399	2,275
R^2	.015	.038	.007	.054
Year Fixed Effects	Yes	Yes	Yes	Yes

Notes: Robust standard errors are shown in parentheses. Each cell presents a regression discontinuity (RD) estimator using household-level data in panel A and individual-level data in panel B. Results are reported using a first-order polynomial control function. In panel A, the dependent variables in columns 1 and 2 are dummy variable indicators for the probability of birth at the household level, nine months after the law was first discussed in the National Assembly in March 2013. The dependent variables in columns 3 and 4 are dummy variable indicators for the probability of birth at the household level, nine months after the law was implemented in April 2014. In panel B, the dependent variables in columns 1 and 2 refer to the number of hours of work per week for women, while the dependent variables in columns 3 and 4 refer to the number of hours of work per week for men. "Half benefits" is a dummy variable indicator equal to 1 for households that receive half the amount of PAJE benefits and equal to 0 for households that receive the full benefits. "Zero benefits" is a dummy variable indicator equal to 1 for households that receive zero PAJE benefits and equal to 0 for households that are eligible for full or half the PAJE benefits. Regressions also include a linear control function, the interaction term between the control function and the RD estimator, and a year fixed effect.

† $p<.10$; * $p<.05$; ** $p<.01$

variable is a dummy variable indicator for births that occurred nine months after the law was implemented. The former definition is my benchmark definition, as it likely provides an upper bound on when any birth effect could possibly occur.[23]

The RDD results suggest that receiving half the amount of the basic allowances of early childhood benefits relative to the full amount did not have any impact on fertility. On the other hand, the total elimination of the allowances decreased the probability of having an (additional) child. The results in Table 2 show that the announcement of the policy reform had a much larger impact on fertility than the actual implementation of the reform. Indeed, being eligible for zero benefits led to a decrease in the birth probability at the household level of 2 percentage points following the announcement of the policy and of 0.7 percentage points following its actual implementation. To compute the implied benefit elasticity, I evaluate these effects relative to the average birth probability at the household level, over a five-year window preceding the reform, for households with women of childbearing age.[24]

Relative to the pre-reform mean, the total elimination of the allowances led to a decline in the probability of having an (additional) child of 38% following the announcement of the reform and of 13% following the policy implementation. These birth effects can be interpreted in terms of benefit elasticities. The elasticity measures the percentage change in the birth probability generated by a 1% change in benefits. The birth effect is observed at the zero benefits threshold, which corresponds to a 100% reduction in benefits. The announcement effect therefore corresponds to an implied elasticity of fertility to benefits of 0.38, while the implementation effect corresponds to an implied elasticity of 0.13.

The results in Table 2 suggest that the largest drop in fertility occurred in 2014, following the earliest discussions of the policy reform, while the actual implementation of the reform led to a lower decline in fertility. There are several explanations for this. For one, as highlighted earlier, the intense media coverage of the policy reform when it was still in process could explain these large short-term impacts. This is also supported by the Google trends analysis, which confirms the surge in web search popularity of the PAJE benefits following the earliest discussions of the reforms. Moreover, aggregate fertility trends in France also support this sharp short-term decline in 2014. Figure 3 plots data from the World Populations Prospects of the United Nations (2019). Indeed, panel b shows a large drop in the annual percentage change in births in the year 2014. While the trend in births was overall declining between 2010 and 2018, as shown in panel a, panel b shows a clear shift in the annual percentage change in births from a pre-reform trend of −0.76% to a post-reform trend of −1.8%.

[23] Relying on the announcement versus implementation definitions of fertility does not impact the treatment and control groups, however, the definitions provide upper and lower bounds for the fertility impact. Indeed, the announcement effect captures the upper bound of when any birth effect could possibly occur, while the implementation effect captures the lower bound. For the latter, there is no change in the definition of treatment and control groups, but if we believe that the announcement effect is the true event, the implementation effect artificially attributes a value of zero for births that occurred nine months after the announcement, leading to a lower bound estimate of the true impact. Contrasting both results therefore provides lower and upper bounds estimation of the fertility response.

[24] The average birth probability at the household level between 2009 and 2013 for households with women of childbearing age (15-49 years old) is .053 using the SRCV dataset.

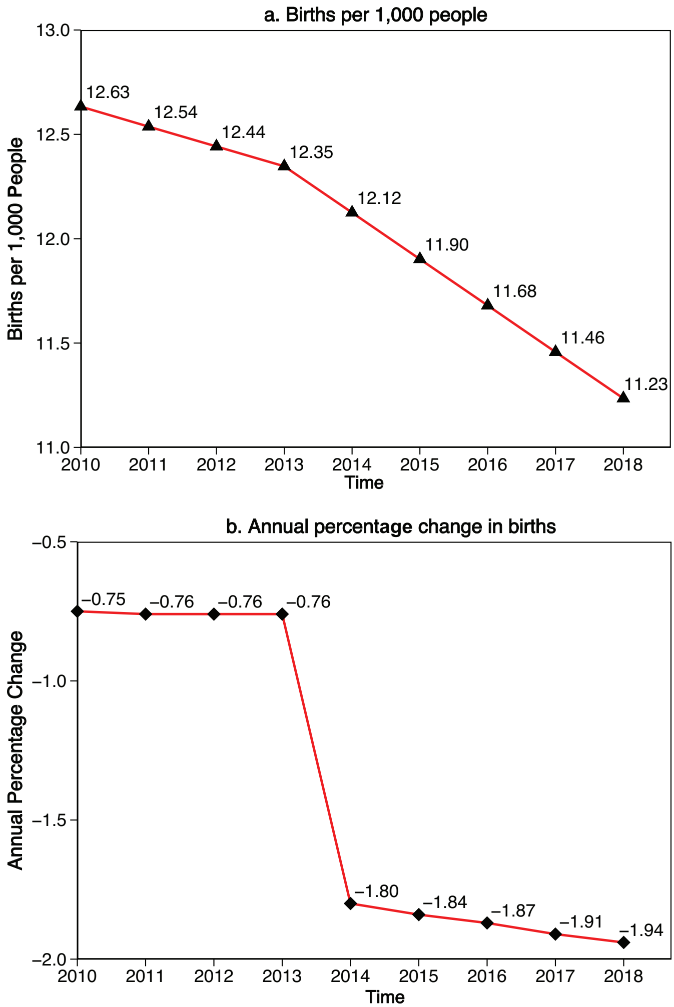

Fig. 3 Aggregate fertility in France before and after the reform. Data come from the World Population Prospects of the United Nations (2019) and the French National Institute of Statistics and Economic Studies (INSEE). Panel a presents the number of births per 1,000 people from 2010 to 2018, while panel b presents the annual percentage change in births (year to year change). Panel c presents the number of monthly births per 1,000 people, as deviations from the monthly mean. A value of 0 on the *y*-axis means that the number of live births for a particular observation is equal to the month mean during the period under consideration, a value greater than 0 means that the number of live births is greater than the month mean, and a value lower than 0 means that the number of live births is lower than the month mean.

(continued)

Furthermore, using monthly birth data from INSEE in panel c, I plot deviations in the number of monthly births (per 1,000 people) from the mean of the corresponding month during the time period under consideration, to address birth seasonality.[25]

[25] The raw monthly data on live births in France showed clear birth seasonality with a peak of births during the month of July and consistently higher number of births in between the months of July and October of each year.

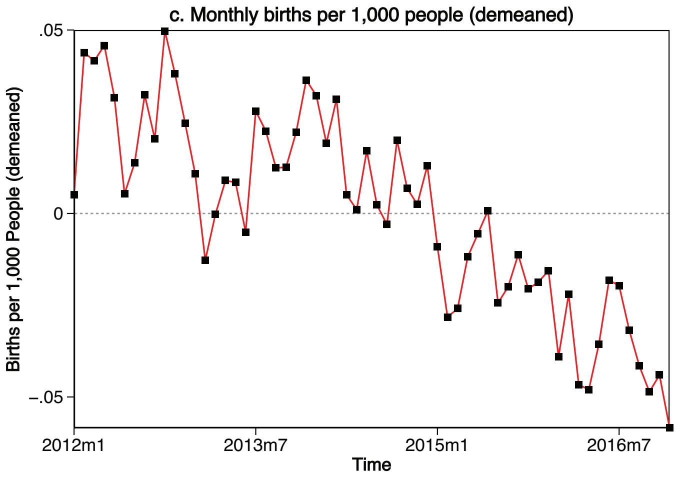

Fig. 3 (continued)

Prior to March 2013, the number of monthly live births was higher than the mean for a given month (values on the *y*-axis greater than 0). This chart also shows a consistent decline thereafter in the number of monthly live births in France, which corresponds to a shift in the trend from higher than average monthly births to lower than average monthly births (values below the horizontal dotted line). The decline in the number of live births is evident nine months after the announcement of the reform, with persistent effects after implementation of the law.

Turning to the results on labor supply in panel B of Table 2, I explore the impact of the 2014 reform on the hours of work per week. The findings suggest that decreasing the amount of basic allowances of early childhood benefits, by receiving either half or zero benefits, was associated with higher labor supply for both women and men. Indeed, receiving half benefits led to an increase in the number of hours of work per week for women by four hours compared with women who belonged to households receiving full benefits. Meanwhile, being eligible for zero benefits increased the number of women's hours of work per week by two compared with women who belonged to households who were eligible for half or the full amount of basic allowances. Relative to the average weekly working hours over a five-year window prior to the reform, the increase relative to the mean was on the order of 10% for women who received half benefits and 5% for those who received zero benefits.[26] The results also suggest that men who received half or zero benefits significantly increased their number of hours of work per week, by two hours and four hours, respectively. The magnitude of these increases relative to the average number of hours of work per week for men over a five-year window prior to the reform was on the order of 6% and 10%, respectively.[27]

[26] Using the SRCV data, the average number of hours of work per week for women of working age (15-64 years old) between 2009 and 2013 was 34.26, while the corresponding average number of hours of work per week for men was 41.11.

[27] In Table A9 in appendix section E1, I also explore heterogeneous effects by nationality. In panel A, the analysis is restricted to French nationals (French by birth, French by naturalization or marriage, or filing

Nonparametric Bandwidth Estimation

In addition to using a control function specification, in Table 3 I also rely on nonparametric bandwidth estimations for households within various income windows from the discontinuity thresholds defined by the policy (±10,000 euros, ±7,500 euros, ±5,000 euros, and ±2,500 euros). The distance from the discontinuity threshold is reported in the last row. In panels A and B, results are reported for birth probability at the household level, while panels C and D report results on the number of hours of work per week for women and men, respectively. In the first two panels, for households within ±10,000 euros or ±7,500 of yearly household income from the discontinuity thresholds defined by the policy, the bandwidth estimation results are fully consistent with the main findings presented in Table 2, using a first-order polynomial control function, and suggest that being eligible for zero benefits is associated with lower birth probability at the household level, relative to households that are eligible for early childhood benefits. When narrowing the income windows to ±5,000 or ±2,500 euros, the results are noisier, as the number of observations decreases and the probability of birth within many income bins is equal to zero. The estimated effect, however, remains negative, although imprecisely estimated.

Panels C and D of Table 3 show that women who are eligible for half benefits work approximately two additional hours per week compared with women who are eligible for the full amount, using a bandwidth of ±10,000 euros or ±7,500. Using narrower bandwidths shows that women who belong to households that are not eligible for any basic allowances of early childhood benefits work two additional hours per week compared with women who are eligible for half benefits. Consistently across all bandwidths, the results in panel D also suggest that being eligible for half or zero benefits is associated with higher working hours for men.

These findings are robust to various checks.[28] In appendix section D1, I present manipulation tests following McCrary (2008) and show that there is no bunching on either side of the cutoff or evidence of manipulation of the forcing variable. In appendix section E2, I provide several robustness checks on sample selection including the exclusion of one-parent households or households that witnessed changes in their composition between the baseline and the time of the survey. In appendix section

after the age of 18), while in panel B, the analysis is restricted to other nationalities such as EU citizens from the countries that entered the EU after 2004; EU citizens from other European countries; Algerian, Moroccan, or Tunisian; African nationals except for the Maghreb countries; and other nationalities or stateless. The results are driven by French nationals, while I do not find any significant impact of the 2014 family policy reforms on fertility responses and labor supply for all other nationalities.

[28] Appendix section B discusses all concomitant reforms (the birth premium, the parental leave reform, the family tax quotient, and the family allowances' reform) and provides several robustness checks suggesting that the results are indeed driven by the withdrawal of the cash benefit and not the other concomitant reforms. The family allowances' reform, which came into force in mid-July 2015, also introduced means-tested family allowances. However, the income thresholds introduced by the reform, as presented in Table A4 in appendix section C, depend on the number of children but are unconditional on household types (one-parent with income, two-parent with one income, or two-parent with two incomes). Moreover, as shown in appendix section C, the family allowances' reform introduced income thresholds that are quite different from those applicable to the PAJE benefits (Table A3). Indeed, the exact thresholds used for the analysis of the PAJE benefits were not employed in any other assistance program that was part of the 2014 or 2015 reforms.

Table 3 Nonparametric bandwidth estimation: Impact of the PAJE reform on fertility and labor supply

	A: Birth Probability (announcement effect)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Half Benefits	–.004		–.010		–.013		.006	
	(.009)		(.010)		(.012)		(.019)	
Zero Benefits		–.014†		–.013†		–.008		–.013
		(.007)		(.008)		(.009)		(.015)
Number of Observations	1,374	1,268	1,158	964	777	649	373	336
R^2	.000	.002	.001	.003	.001	.001	.000	.002
	B: Birth Probability (implementation effect)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Half Benefits	.006		.004		.003		.000	
	(.004)		(.004)		(.006)		(.011)	
Zero Benefits		–.008*		–.009†		–.006		–.012
		(.004)		(.005)		(.004)		(.009)
Number of Observations	1,374	1,268	1,158	964	777	649	373	336
R^2	.002	.032	.001	.003	.000	.003	.000	.006
	C: Hours of Work/Week for Women							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Half Benefits	1.256*		1.275*		1.165		0.859	
	(0.533)		(0.582)		(0.727)		(0.919)	
Zero Benefits		0.310		0.674		1.576*		2.088*
		(0.564)		(0.641)		(0.769)		(0.975)
Number of Observations	1,167	1,106	998	857	656	588	304	310
R^2	.005	.000	.005	.001	.004	.007	.003	.015
	D: Hours of Work/Week for Men							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Half Benefits	1.458**		1.513*		1.889*		1.115	
	(0.551)		(0.603)		(0.733)		(1.026)	
Zero Benefits		1.707**		1.418*		1.617†		2.116†
		(0.605)		(0.693)		(0.852)		(1.112)
Number of Observations	1,248	1,230	1,069	942	732	637	349	324
R^2	.006	.007	.006	.005	.009	.006	.003	.011
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Distance From Threshold	±10,000	±10,000	±7,500	±7,500	±5,000	±5,000	±2,500	±2,500

Notes: Robust standard errors are shown in parentheses. Each cell presents a regression discontinuity estimator using household-level data in panels A and B and individual-level data in panels C and D. Analysis is restricted to households within ±10,000 euros of yearly household income from the income threshold defined by the 2014 policy reform in columns 1 and 2, within ±7,500 euros in columns 3 and 4, within ±5,000 euros in columns 5 and 6, and within ±2,500 euros in columns 7 and 8. In panel A, the dependent variable is a dummy variable indicator for the probability of birth at the household level, nine months after the law was first discussed in March 2013. In panel B, the dependent variable is a dummy variable indicator for the probability of birth at the household level, nine months after the law came into force in April 2014. In panels C and D, the dependent variables refer to the number of hours of work per week for women and men, respectively. “Half benefits” is a dummy variable indicator equal to 1 for households that receive half the amount of PAJE benefits and to 0 for households that receive the full PAJE benefits. “Zero benefits” is a dummy variable indicator equal to 1 for households that receive zero PAJE benefits and to 0 for households that are eligible for full or half the PAJE benefits. Distance from the discontinuity threshold is reported in the last row.

†$p<.10$; *$p<.05$; **$p<.01$

Table 4 Heterogeneous effects by birth parity: Impact of the PAJE reform on fertility

	Announcement Effect				Implementation Effect			
	First-born (1)	Parity 2+ (2)	First-born (3)	Parity 2+ (4)	First-born (5)	Parity 2+ (6)	First-born (7)	Parity 2+ (8)
Half Benefits	–.005	.001			.000	.006		
	(.003)	(.018)			(.000)	(.009)		
Zero Benefits			–.002	–.018*			.000	–.007†
			(.002)	(.008)			(.000)	(.004)
Number of Observations	1,890	1,890	2,786	2,786	1,890	1,890	2,786	2,786
R^2	.011	.020	.008	.019	.000	.019	.000	.015
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Robust standard errors are shown in parentheses. Each cell presents a regression discontinuity (RD) estimator using household-level data and a first-order polynomial control function. The dependent variables are dummy variable indicators for the probability of birth at the household level, for births that occurred nine months after the law was first discussed in the National Assembly in March 2013. "Half benefits" is a dummy variable indicator equal to 1 for households that receive half the amount of PAJE benefits and equal to 0 for households that receive the full PAJE benefits. "Zero benefits" is a dummy variable indicator equal to 1 for households that receive zero PAJE benefits and equal to 0 for households that are eligible for full or half the PAJE benefits. Regressions also include a linear control function, the interaction term between the control function and the RD estimator, and a year fixed effect.

†$p < .10$; *$p < .05$

E3, I report results using a weighted first-order polynomial control function specification, including a vector of predetermined individual and household covariates, and using second-order and third-order specifications of the control function to account for nonlinearities in income. Appendix section E4 also shows that the results are robust to dropping all heaped income values and to alternative definitions of household income and of the working/active status of individuals. In appendix section E5, I show that the results are robust to using standard panel data analysis combined with a difference-in-differences approach. Finally, the placebo regressions presented in appendix section E6—using pre-reform data or relying on placebo thresholds—suggest that there are no preexisting trends in fertility and labor supply, while appendix section E7 provides evidence that fertility-related attrition is not driving my results.

Heterogeneity Analysis and Underlying Mechanisms

In Table 4, I also examine the effects of the PAJE reform on fertility by disentangling the effect on the probability of becoming parents for households without children and the probability of having an additional child for current parents. To do so, I differentiate between first-born children and children of higher birth parity (2+). Indeed, the results suggest that the policy reform only reduced the probability of having a second child or a child of higher birth parity without affecting the decision of becoming parents for households without children. In line with the benchmark results reported in

Table 5 Underlying mechanisms: Impact of the PAJE reform on labor supply

	A: Switching From Part-time to Full-time			
	Women		Men	
	(1)	(2)	(3)	(4)
Half Benefits	.005		.004	
	(.016)		(.009)	
Zero Benefits		.006		–.001
		(.008)		(.004)
Number of Observations	1,273	2,022	1,295	2,154
R^2	.003	.002	.002	.003
	B: Having Multiple Jobs			
	Women		Men	
	(1)	(2)	(3)	(4)
Half Benefits	.002		.010	
	(.020)		(.015)	
Zero Benefits		–.002		.000
		(.010)		(.009)
Number of Observations	1,409	2,198	1,407	2,283
R^2	.007	.011	.001	.000
	C: Labor Market Participation			
	Women		Men	
	(1)	(2)	(3)	(4)
Half Benefits	.011		–.028	
	(.041)		(.038)	
Zero Benefits		.013		–.005
		(.023)		(.017)
Number of Observations	2,402	3,567	2,072	3,275
R^2	.005	.007	.006	.001

(continued)

Table 2, the announcement of the policy had a much larger impact on fertility relative to its actual enactment. Relative to the pre-reform mean, the total elimination of the allowances led to a decline in the probability of having a second child or a child of higher birth parity by 34% following the announcement of the reform and by 13% following the policy implementation.

I also explore several underlying mechanisms for the labor supply response. Table A10 presents heterogeneous effects by type of employment (full-time vs. part-time) and suggests that the increase in labor supply is entirely driven by full-time employment. To further investigate the underlying mechanisms, in Table 5, I consider several competing channels. First, I examine whether there is an increase in the probability of switching from part-time to full-time employment around the discontinuity thresholds. Indeed, the SRCV dataset provides information on whether an individual is employed on a full-time or part-time basis in April of the year considered, but also in April of the previous year. Second, I explore whether the increase in labor supply could be

Table 5 (continued)

	D: Overtime (>35 hours/week)			
	Women		Men	
	(1)	(2)	(3)	(4)
Half Benefits	.136*		.119**	
	(.056)		(.036)	
Zero Benefits		.125**		.118**
		(.035)		(.029)
Number of Observations	1,405	2,194	1,399	2,275
R^2	.010	.040	.010	.035
Year Fixed Effects	Yes	Yes	Yes	Yes

Notes: Robust standard errors are shown in parentheses. Each cell presents a regression discontinuity (RD) estimator using individual-level data. Results are reported using a first-order polynomial control function. In panel A, the dependent variable is a dummy variable indicator equal to 1 for individuals who switch from part-time employment to full-time employment (between April of the year before and April of the year considered). In panel B, the dependent variable is a dummy variable indicator equal to 1 for individuals who reported having multiple jobs. In panel C, the dependent variable is a dummy variable indicator for labor market participation. In panel D, the dependent variable is a dummy variable indicator equal to 1 for individuals working more than 35 hours per week, which corresponds to the legal number of working hours. "Half benefits" is a dummy variable indicator equal to 1 for households that receive half the amount of PAJE benefits and equal to 0 for households that receive the full PAJE benefits. "Zero benefits" is a dummy variable indicator equal to 1 for households that receive zero PAJE benefits and equal to 0 for households that are eligible for full or half the PAJE benefits. Regressions also include a linear control function, the interaction term between the control function and the RD estimator, and a year fixed effect.

*$p<.05$; **$p<.01$

due to having multiple jobs. Third, I look at the impact on labor market participation to see whether the increase in the number of weekly hours of work could be due to individuals reentering the labor market earlier. Finally, I explore whether the increase in labor supply could be attributed to overtime (working more than the legal duration of 35 hours per week). This latter possibility relies on the definition of the Research Division of the French Ministry of Labor (Direction de l'Animation de la Recherche, des Études et des Statistiques, DARES), which states that overtime hours correspond to the hours worked beyond the legal working time set at 35 hours per week.[29]

The Table 5 findings suggest that the increase in labor supply cannot be explained by any one of the first three channels and rather support that the increase in labor supply is due to working overtime. Indeed, we found that women and men who experienced a reduction in the amount of PAJE benefits had a greater likelihood of working overtime (more than 35 hours per week). For example, receiving half the amount of basic allowances of early childhood benefits or not receiving any was associated with a 13% increase in the probability of working overtime for women and an 11% increase in overtime for men. Appendix section F provides supplementary

[29] See "Les heures supplémentaires rémunérées," Direction de l'Animation de la Recherche, des Études et des Statistiques (DARES), https://dares.travail-emploi.gouv.fr/donnees/les-heures-supplementaires-remunerees.

official data on overtime among public- and private-sector employees, published by French government institutions, supporting the underlying mechanism for the labor supply response.

Conclusions

This study examines the impact of the 2014 family policy reforms of PAJE benefits in France on fertility choices and labor supply for both women and men. The reform provides a quasi-experimental setting as it made the basic allowances of early childhood benefits a piece-wise constant function of families' past income. The reform defined income thresholds beyond which households became either eligible for half the amount of basic allowances or ineligible for the basic allowances of early childhood benefits. Exploiting this "sharp" discontinuity in the provision of early childhood benefits, I examine its impact on birth probability at the household level and on women's and men's labor supply.

The analysis relies on an RDD utilizing data from the European Union–Statistics on Income and Living Conditions in France for the years 2014 and 2015. One of the main advantages of the RDD framework is the local quasi-randomization around the specific thresholds employed in the analysis, which ensures that households close to the discontinuity threshold are very comparable along observable and unobservable characteristics and, by extension, that the results estimated in the neighborhood of the discontinuity thresholds solely capture the effects of the reform rather than the effects of other economic factors—which are, unlike the reform in question, not randomized around the discontinuity thresholds.

The presented results suggest that halving the amount of basic allowances does not have any impact on fertility; however, the total elimination of the benefits is found to be associated with lower fertility. The study captures both an "announcement effect" and an "implementation effect" of the policy reform: the largest decline in fertility occurred following the earliest discussion of the policy reform (an elasticity of 0.38). Following its implementation, the estimated benefit elasticity was 0.13. The larger fertility decline following the announcement is likely due to the extensive media coverage of the reform process and might be suggestive of timing effects.

To further understand these dynamics, I explore whether the decline in fertility reflects a decline in the probability of having children for households without children prior to the reform or a decline in the probability of having an additional child. The results suggest that the PAJE benefits reform affected the decision of having additional children but not the decision of becoming parents for the first time. Another important question is whether the drop in fertility associated with the total elimination of the allowances reflects actual decreases in fertility (a "quantum effect") or merely changes in the timing of births (a "tempo effect"). Given the short time span between the 2014 policy reform and the time of the surveys employed in this study, it is challenging to clearly identify and disentangle these two mechanisms. Indeed, one would need to study fertility over a longer period of time to assess whether the decline in fertility associated with the total elimination of the allowances reflects a decline in completed fertility. Nonetheless, this would not be possible since households could manipulate their income information to alter eligibility.

Regarding the impact of the 2014 family policy reforms on labor supply, I find consistent results in line with the literature on the elimination of welfare programs. As posited by standard labor supply theory, the reduction in child benefits—for households that became eligible for either half the amount of PAJE benefits or households that became ineligible for any—is associated with an increase in the number of hours of work per week for both women and men. A back-of-the-envelope calculation shows that the implied change in earned income, due to an increase in weekly working hours, corresponds with the euro value reduction in benefits.

These results are robust to a battery of robustness checks and, importantly, the online appendix provides evidence that the findings are not driven by any other concurrent reform in France, including the birth premium reform, the family tax quotient cap, the parental leave reform, and the family allowances reform.

The estimated effects presented here reflect a local average treatment effect due to one-sided noncompliance. The noncompliers are those who would not receive the treatment despite eligibility (the never-takers). As discussed thoroughly in appendix section G, these households or individuals are likely to belong to the top income deciles, as shown in recent statistics from French Social Security. Given that richer households have lower fertility but also consistently higher labor supply, the fertility estimates provided likely reflect a lower bound of the true effect, while the labor supply response provides an upper bound.

These results carry important policy implications, not only for France but also for other OECD countries. Indeed, the decline in fertility and the effect of cash benefits constitute ongoing policy concerns in many OECD countries. While the empirical findings are specific to France, the policy implications are far-reaching. In recent years, public spending on cash family benefits in OECD countries dropped from 1.4% of GDP in 2009 to 1.2% of GDP in 2017 (OECD 2018). Over the same period, fertility within OECD countries declined from 1.8 to 1.7 births per woman (World Bank 2020). While fertility in OECD countries has been consistently declining since 1960, reaching its lowest level in 2020 at 1.6 births per woman (World Bank 2020), and the reduction in cash benefits is but one of many factors that could impact these trends, the evidence from France suggests that cash benefits matter in determining households' fertility and labor market decisions. These findings suggest that cash benefit withdrawal or reductions might impact fertility in other contexts. Further research would be needed to examine whether such dynamics apply in other countries. ■

Acknowledgments This work was supported by the ANR FamPol Project (grant ANR-14-FRAL-0007) of the French Agence Nationale de la Recherche. I would like to thank the *Demography* editor in chief, Sara R. Curran, the deputy editor, and two anonymous referees for their valuable feedback. The author also thanks Francesco D'Amuri, Michael Darden, Marta De Philippis, Nicolas Jacquemet, Lawrence Katz, Bertrand Koebel, Stefano Neri, Matthew Notowidigdo, Giovanni Peri, Robert M. Sauer, Biagio Speciale, Andrea Weber, and Ekaterina Zhuravskaya for useful comments.

References

Adda, J., Dustmann, C., & Stevens, K. (2017). The career costs of children. *Journal of Political Economy, 125,* 293–337.

Albrecht, J. W., Edin, P.-A., Sundström, M., & Vroman, S. B. (1999). Career interruptions and subsequent earnings: A reexamination using Swedish data. *Journal of Human Resources, 34,* 294–311.

Allocations familiales: Les propositions du rapport Fragonard [Family allowances: The proposals of the Fragonard report]. (2013, April 1). *Le Monde*. Retrieved from https://www.lemonde.fr/politique/article/2013/04/01/allocations-familiales-les-propositions-du-rapport-fragonard_3151644_823448.html

Arnault, S. (2018). *Salaire horaire: L'importance de la catégorie socio-professionnelle et du diplôme* [Hourly wage: The importance of socio-professional category and diploma] (INSEE Focus Report, No. 116). Retrieved from https://www.insee.fr/fr/statistiques/3564663#consulter

Averett, S. L., & Whittington, L. A. (2001). Does maternity leave induce births? *Southern Economic Journal, 68,* 403–417.

Baker, M., & Milligan, K. (2008). Maternal employment, breastfeeding, and health: Evidence from maternity leave mandates. *Journal of Health Economics, 27,* 871–887.

Bauernschuster, S., Hener, T., & Rainer, H. (2016). Children of a (policy) revolution: The introduction of universal child care and its effect on fertility. *Journal of the European Economic Association, 14,* 975–1005.

Baum, C. L. (2003). The effect of state maternity leave legislation and the 1993 Family and Medical Leave Act on employment and wages. *Labour Economics, 10,* 573–596.

Becker, G. S. (1960). An economic analysis of fertility. In Universities–National Bureau Committee for Economic Research (Ed.), *Demographic and economic change in developed countries* (pp. 209–240). New York, NY: Columbia University Press. Retrieved from https://www.nber.org/system/files/chapters/c2387/c2387.pdf

Berger, L. M., & Waldfogel, J. (2004). Maternity leave and the employment of new mothers in the United States. *Journal of Population Economics, 17,* 331– 349.

Björklund, A. (2006). Does family policy affect fertility? *Journal of Population Economics, 19,* 3–24.

Brewer, M., Ratcliffe, A., & Smith, S. (2012). Does welfare reform affect fertility? Evidence from the UK. *Journal of Population Economics, 25,* 245–266.

Brunner, B., & Kuhn, A. (2014). Announcement effects of health policy reforms: Evidence from the abolition of Austria's baby bonus. *European Journal of Health Economics, 15,* 373–388.

Card, D., & Lemieux, T. (1997). *Adapting to circumstances: The evolution of work, school, and living arrangements among North American youth* (NBER Working Paper 6142). Cambridge, MA: National Bureau of Economic Research.

Cohen, A., Dehejia, R., & Romanov, D. (2013). Financial incentives and fertility. *Review of Economics and Statistics, 95,* 1–20.

Crump, R., Goda, G. S., & Mumford, K. J. (2011). Fertility and the personal exemption: Comment. *American Economic Review, 101,* 1616–1628.

Dustmann, C., & Schönberg, U. (2012). Expansions in maternity leave coverage and children's long-term outcomes. *American Economic Journal: Applied Economics, 4*(3), 190–224.

Gans, J. S., & Leigh, A. (2009). Born on the first of July: An (un)natural experiment in birth timing. *Journal of Public Economics, 93,* 246–263.

Gauthier, A. H., & Hatzius, J. (1997). Family benefits and fertility: An econometric analysis. *Population Studies, 51,* 295–306.

González, L. (2013). The effect of a universal child benefit on conceptions, abortions, and early maternal labor supply. *American Economic Journal: Economic Policy, 5*(3), 160–188.

González, L., & Trommlerová, S. K. (2023). Cash transfers and fertility: How the introduction and cancellation of a child benefit affected births and abortions. *Journal of Human Resources, 58,* 783–818.

Guichard, G. (2013, June 4). Baisse du quotient familial: Ce que cela va coûter aux familles [Declining family quotient: What it will cost families]. *Le Figaro*. Retrieved from https://www.lefigaro.fr/conjoncture/2013/06/03/20002-20130603ARTFIG00670-baisse-du-quotient-familial-ce-que-cela-va-couter-aux-familles.php#:~:text=Elle%20atteindra%20500%20euros%20maximum,soit%2021%20euros%20par%20mois

Haan, P., & Wrohlich, K. (2011). Can child care policy encourage employment and fertility?: Evidence from a structural model. *Labour Economics, 18,* 498–512.

Hoem, J. M. (1993). Public policy as the fuel of fertility: Effects of a policy reform on the pace of childbearing in Sweden in the 1980s. *Acta Sociologica, 36,* 19–31.

Hotz, V. J., Klerman, J. A., & Willis, R. J. (1997). The economics of fertility in developed countries. In M. R. Rosensweig & O. Stark (Eds.), *Handbooks in economics: Vol. 14. Handbook of population and family economics* (Vol. 1A, pp. 275–347). Amsterdam, the Netherlands: North-Holland.

Joyce, T., Kaestner, R., Korenman, S., & Henshaw, S. (2004). Family cap provisions and changes in births and abortions. *Population Research and Policy Review, 23,* 475–511.

Kalwij, A. (2010). The impact of family policy expenditure on fertility in western Europe. *Demography, 47,* 503–519.

Kearney, M. S. (2004). Is there an effect of incremental welfare benefits on fertility behavior? A look at the family cap. *Journal of Human Resources, 39,* 295– 325.

Kim, Y.-I. A. (2014). Lifetime impact of cash transfer on fertility. *Canadian Studies in Population, 41,* 97–110.

Lalive, R., & Zweimüller, J. (2009). How does parental leave affect fertility and return to work? Evidence from two natural experiments. *Quarterly Journal of Economics, 124,* 1363–1402.

Laroque, G., & Salanié, B. (2014). Identifying the response of fertility to financial incentives. *Journal of Applied Econometrics, 29,* 314–332.

Laurent, S., & Laurent, S. (2013, April 9). Prestations familiales: De quoi parle-t-on? [Family benefits: What are we talking about?]. *Le Monde.* Retrieved from https://www.lemonde.fr/decryptages/article/2013/04/09/prestations-familiales-pourquoi-une-reforme_3151865_1668393.html

Le gouvernement choisit d'abaisser le plafond du quotient familial [The government chooses to lower the ceiling of the family quotient]. (2013, June 3). *Le Figaro*. Retrieved from https://www.lefigaro.fr/conjoncture/2013/06/03/20002-20130603ARTFIG00425-le-gouvernement-choisit-d-abaisser-le-plafond-du-quotient-familial.php#:~:text=C%27est%20la%20solution%20d,euros%20par%20demi%2Dpart%20fiscale

Lemieux, T., & Milligan, K. (2008). Incentive effects of social assistance: A regression discontinuity approach. *Journal of Econometrics, 142,* 807–828.

Malkova, O. (2018). Can maternity benefits have long-term effects on childbearing? Evidence from Soviet Russia. *Review of Economics and Statistics, 100,* 691–703.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics, 142,* 698–714.

Milligan, K. (2005). Subsidizing the stork: New evidence on tax incentives and fertility. *Review of Economics and Statistics, 87,* 539–555.

Moffitt, R. A. (2002). Welfare programs and labor supply. In A. J. Auerbach & M. Feldstein (Eds.), *Handbook of public economics* (Vol. 4, pp. 2393–2430). Amsterdam, the Netherlands: Elsevier Science.

OECD. (2018). *Family benefits public spending* (Indicator) [Dataset]. Organisation for Economic Cooperation and Development. https://doi.org/10.1787/8e8b3273-en

OECD. (2021). *Enrollment in childcare and pre-school* [Dataset]. Organisation for Economic Cooperation and Development. Retrieved from https://www.oecd.org/els/soc/PF3_2_Enrolment_childcare_preschool.pdf

Olsen, R. J. (1994). Fertility and the size of the U.S. labor force. *Journal of Economic Literature, 32,* 60–100.

Parent, D., & Wang, L. (2007). Tax incentives and fertility in Canada: Quantum vs tempo effects. *Canadian Journal of Economics / Revue Canadienne d'Economique, 40,* 371–400.

Pettersson-Lidbom, P. (2008). Do parties matter for economic outcomes? A regression–discontinuity approach. *Journal of the European Economic Association, 6,* 1037–1056.

Piketty, T. (2005). Impact de l'allocation parentale d'éducation sur l'activité féminine et la fécondité en France [Impact of the parental education allowance on female activity and fertility in France]. In C. Lefèvre & A. Filhon (Eds.), *Les cahiers de l'Ined, Vol. 156. Histoires de familles, histoires familiales: Les résultats de l'Enquête Famille de 1999* [Family stories, family histories: The results of the 1999 Family Survey] (pp. 79–109). Paris, France: Ined Éditions.

Raute, A. (2019). Can financial incentives reduce the baby gap? Evidence from a reform in maternity leave benefits. *Journal of Public Economics, 169,* 203–222.

Riphahn, R. T., & Wiynck, F. (2017). Fertility effects of child benefits. *Journal of Population Economics, 30,* 1135–1184.

Rosenzweig, M. R. (1999). Welfare, marital prospects, and nonmarital childbearing. *Journal of Political Economy, 107*(S6), S3–S32.
Ruhm, C. J. (1998). The economic consequences of parental leave mandates: Lessons from Europe. *Quarterly Journal of Economics, 113,* 285–317.
Schøne, P. (2004). Labour supply effects of a cash-for-care subsidy. *Journal of Population Economics, 17,* 703–727.
United Nations. (2019). *World population prospects 2019* (Online ed., revision 1) [Dataset]. United Nations, Department of Economic and Social Affairs, Population Division. Retrieved from https://population.un.org/wpp2019/
Vignaud, M. (2013, June 3). Quotient familial: Le gouvernement augmente l'impôt des familles les plus aisées [Family quotient: The government increases taxes for the wealthiest families]. *Le Point.* Retrieved from https://www.lepoint.fr/economie/le-gouvernement-augmente-l-impot-des-familles-les-plus-aisees-03-06-2013-1675944_28.php#11
Whittington, L. A. (1992). Taxes and the family: The impact of the tax exemption for dependents on marital fertility. *Demography, 29,* 215–226.
Whittington, L. A., Alm, J., & Peters, H. E. (1990). Fertility and the personal exemption: Implicit pronatalist policy in the United States. *American Economic Review, 80,* 545–556.
World Bank. (2020). *World development indicators: Fertility rate, total (births per woman)* [Dataset]. Retrieved from https://data.worldbank.org/indicator/SP.DYN.TFRT.IN

Nelly Elmallakh
nelmallakh@worldbank.org

World Bank, Washington, DC, USA; https://orcid.org/0000-0003-1011-1479