

Income effects and labour supply: Evidence from a child benefits reform[☆]Mathias Fjællegaard Jensen^{a,*}, Jack Blundell^b^a Department of Economics, University of Oxford, 10 Manor Rd, Oxford, OX1 3UQ, UK^b Centre for Economic Performance, London School of Economics, UK

ARTICLE INFO

JEL classification:

H24

J13

J16

J18

J22

Keywords:

Labour supply

Gender differences

Child benefits

ABSTRACT

In this paper, we exploit a unique and unexpected reform to the child benefit system in Denmark to assess the effects of child benefits on parental labour supply. A cap on child benefit payments in 2011 led to a non-negligible reduction in child benefits for larger families with young children while leaving child benefits for smaller families unchanged. The differential impact of this policy represents an opportunity to assess the causal impact of child benefit programmes on the labour supply of mothers and fathers. Using a difference-in-differences strategy, we find that the reduction in benefits leads to a substantial increase in the labour supply of mothers. Mothers respond to the policy at both the intensive and extensive margins, with the latter outweighing the former, and the effect persists after controlling for fertility-related family characteristics. To fix preferences for additional children across treatment and control groups, we use data on parents' medical consultations on sterilisation, a common procedure in Denmark.

1. Introduction

Financial assistance to families with young children is a common policy adopted across many developed countries. Child or family benefits are cash transfers to families with dependent children and are often independent of income and labour market status. These unconditional benefits represent an alternative to conditional or in-work benefits, such as the federal Earned Income Tax Credit (EITC) in the United States. Child benefits are a major part of public spending in most European countries, with spending amounting on average to 1.16% of GDP across OECD countries in 2017.¹ There are multiple motivations for government spending on child benefits, such as the wellbeing of families, opportunities for children, and incentive effects on fertility. Recently, the US child tax credit was expanded to include monthly allowances for families with children, and thus, a critical discussion of the effects of unconditional child benefits has reemerged (Corinth et al., 2021; Financial Times, 2021).

Child benefits represent a subsidy to parents that enables the maintenance of income levels while limiting their labour supply. We are particularly interested in this effect. Limiting parental labour supply may be beneficial for certain (child) outcomes, but if experienced asymmetrically across parents, it may also reinforce societal outcomes

that policymakers may find less desirable, for example, the child pay penalty and the gender pay gap (see e.g. Kleven et al., 2019; Blau and Kahn, 2017).

A simple economic analysis would predict that elevated (unconditional) child benefits increase fertility and decrease labour supply among parents. With child benefits, the cost of an additional child is relatively lower, and thus, fertility is expected to increase. At the same time, an increase in unearned income generates income effects. The marginal utility of income decreases, predicting a lower level of labour supply when child benefits are implemented. Despite the prevalence of these policies, relatively little is known about whether these theoretical predictions hold in practice. Most of the current evidence on the effects of child benefits rely on analyses comparing families with and without children, despite the fact that these types of families likely differ along many other unobservable dimensions.

We contribute to this literature by exploiting a unique and unexpected reform to the child benefit system in Denmark to assess the effects on maternal and paternal labour supply. A cap on child benefit payments in 2011 led to a non-negligible reduction in child benefits for larger families with young children, but did not directly affect families with one or two children. The differential impact of this policy shift

[☆] We thank seminar and conference participants at Stanford, Copenhagen Business School, EALE 2021, DGPE, and the 5th Bank of Italy-CEPR Labour Workshop. This paper has benefited substantially from discussions with Fane Groes, Herdis Steingrimsdottir, Miriam Wüst, Thomas Jørgensen, Jakob Søgaard, and Caroline Hoxby. We would also like to thank the editor and three anonymous reviewers for their comments, which have greatly benefited our paper. Any errors are our own.

* Corresponding author.

E-mail address: mathias.jensen@economics.ox.ac.uk (M.F. Jensen).

¹ <https://data.oecd.org/socialexp/family-benefits-public-spending.htm>

represents an opportunity to assess the causal impact of child benefit programmes on the labour supply of mothers and fathers, respectively. The cap on child benefits was not directly related to family income levels, but rather based only on the number and age composition of their children. Given this, all estimated labour supply effects are due to income effects alone. This setting contrasts that of expansion of EITC in the US and other in-work child benefits where both substitution and income effects need to be addressed. In light of this, we can compare families with different compositions of children and leverage population-level Danish employment registers to control for individual fixed effects.

We find that mothers respond to a reduction in child benefits both on the intensive and extensive margins. Mothers increase the number of hours worked per month by 1.17 if employed. On top of this, we find an extensive margin response of a 1.36 percentage point increase in participation rates for mothers. We show that this large-magnitude response is not due to the indirect effects of changes in fertility, but rather directly due to the income effect of the policy shift. This is demonstrated firstly by controlling for fertility-related family characteristics, and secondly by limiting our sample to families who have consulted a doctor regarding sterilisation. We also consider heterogeneity in the response to the reform, finding that mothers with the highest earnings potential respond relatively strongly to the reform.

When considering the combined intensive and extensive margin responses, our results imply that mothers on average increase their labour supply by 2.69 h per month, corresponding to a 2.25% increase in hours worked. Our best estimate of the female labour supply elasticity with respect to income is -0.900 including responses at both the extensive and intensive margins, and -0.387 when only considering the response at the intensive margin. These elasticities are large in absolute values compared to most of those in the existing literature, which often finds values near zero for married and cohabiting mothers (see e.g. Hotz and Scholz, 2003; McClelland and Mok, 2012, for reviews). However, also Eissa and Hoynes (2004) and Eissa and Hoynes (2006b) find statistically significant and negative female labour supply elasticities, as they document a substantial reduction in maternal labour supply with the expansion of the EITC for married couples in the US. Eissa and Hoynes (2006b) estimate an income elasticity between -0.04 and -0.36 at the intensive margin, which is close to our best estimate of -0.387 . We may expect relatively large female labour supply elasticities in absolute values in our setting for a number of reasons. First, we consider families who have finalised fertility, and mothers' labour supply may be more elastic if no further children are expected. Second, the 2011 cap on child benefits considered in this paper is the first substantial deterioration in the provision of child benefits in Denmark. A perceived change in the stability of the child benefit system may induce a "cultural shift" amongst the affected families; that is, a shift away from the culture in which child benefits were perceived as a reliable source of income. Third, income effects may be asymmetric; an income gain and income loss of similar magnitudes may cause different labour supply responses in absolute terms, for example due to credit constraints. Fourth, in Denmark, the labour force participation of women – including mothers – is especially high relative to other countries (OECD, 2020). The generally high labour market participation of mothers in Denmark reflects institutions and norms that support maternal labour supply, which we would expect to influence elasticities.

Although the cap on child benefits was intended to be permanent and declared as such on announcement, a new government was elected in Denmark in November 2011, and the cap was repealed from 2012 onwards. The labour supply effects of the reform can still be seen three years after its repeal. This demonstrates evidence of long-term effects of the ultimately temporary income shock, which can be explained by labour market entry/switching costs and/or by a cultural shift in response to increased uncertainty over future child benefit payments. Because the decline in income from child benefits was ultimately temporary, we find that on net, families' income increased after the

introduction of the reform. The increase in overall income appears to be driven by both mothers and fathers in affected families, which suggests that fathers respond on margins other than hours worked and participation, e.g. by finding a job with a higher wage.

The paper is structured as follows: In the next section, we consider related literature. In Section 3, we describe the Danish child benefits system and the 2011 cap on child benefits in detail. In Section 4, we outline the hypotheses we test in the following sections. In Section 5, we describe our data and define our treatment and control groups. In Section 6, we outline our empirical strategy. In Section 7, we report our estimated effects on aggregate labour supply, and in Section 8 we decompose the responses to the reform into adjustments at the extensive and intensive margins of labour supply. In Section 9, we consider heterogeneity in response to the reform, as well as effects on income and earnings. Finally, in Section 10, we conclude.

2. Literature

Our analysis of the Danish 2011 cap on child benefits contributes to at least two large literatures. Firstly, it relates to the specific literature on the effects of child benefits, family cash transfers, and family tax credits on labour market outcomes, for example, the many papers on the EITC in the US and the 1996-reform of child benefits in Germany. Secondly, our analysis relates to a broader body of studies on income effects. Importantly, the income shock we study in this paper is not directly influenced by labour market status and prior labour market income. In this review, we focus on the former literature, specifically on child benefits. A more general discussion of income effects can be found in, e.g. Hotz and Scholz (2003) and McClelland and Mok (2012).

Labour supply effects of conditional family/child benefits are well-documented. Evaluations of the effects of the US EITC generally show that single mothers' labour supply increases with the introduction and expansions of the tax credit (Eissa and Liebman, 1996; Eissa and Hoynes, 2006a; Gelber and Mitchell, 2012; Kuka and Shenhav, 2020; Meyer and Rosenbaum, 2000; Schanzenbach and Strain, 2021). For example, by comparing single mothers with single women without children, Eissa and Liebman (1996) find that single mothers increased their labour supply at the extensive margin in response to expansions of the EITC in the late 1980s, but they find no effects on the intensive margin. Evaluations of the Working Families' Tax Credit (WFTC) in the UK show similar results, see e.g. Francesconi and van der Klaauw (2007). Thus, existing studies generally find that in-work child subsidies positively affect labour supply of mothers, or at least, single mothers. As an exception in the literature, Kleven (2019) finds little effect of the EITC on employment. Schanzenbach and Strain (2021) discuss the inconsistent findings and emphasise that the results are sensitive to the choice of empirical approach. Using detailed administrative data from California, Hotz and Scholz (2006) again consider the labour supply effects of the EITC. Hotz and Scholz (2006) exploit variation in the EITC by number of children and control for family fixed effects. They also find that expansions of the EITC increase family labour supply. In contrast to the papers on single mothers, Eissa and Hoynes (2004) consider the labour supply responses of married couples to EITC expansions between 1984 and 1996. They find that women decrease their labour supply at the extensive margin and that the labour supply responses of their spouses do not offset this effect, concluding "*the EITC is effectively subsidizing married mothers to stay home...*" (Eissa and Hoynes, 2004, p. 1931). Considering the response at the intensive margin, Eissa and Hoynes (2006b) also find that mothers in married couples significantly reduce their labour supply in response to expansions of the EITC.

In addition to in-work child benefits or tax credits, many countries also pay unconditional child benefits to families with children. Only few studies have evaluated the effects of these – often expensive – policies. For example, Hener (2016) and Tamm (2010) consider the effects of the 1996-increase in child benefit payments in Germany. Using a difference-in-differences setup, Hener (2016) and Tamm

(2010) compare couples with and without children. Tamm (2010) finds that mothers decrease their labour supply on the intensive margin after the increase in child benefits, while fathers' labour supply is unaffected. Hener (2016) observe that the policy's effectiveness in improving families' financial situation is limited while the strain on public finances are amplified by the behavioural response to the increase in child benefits as the resulting decrease in maternal labour supply reduces tax payments. Nevertheless, Raschke (2016) exploits the same reform and shows that families' expenditures on food increase after the increase in child benefits. In addition to the German reform, González (2013) study the introduction of an unconditional child benefit in Spain, namely one-time payment of €2500 after birth. González (2013) find positive fertility effects, no effects on child expenditures, and negative labour supply responses for mothers.

In comparison to the existing papers on the effects of unconditional child benefits, our analysis differs in at least the following ways: (1) We study a cap on child benefit payments that reduced generosity, rather than an increase in generosity or an introduction of child benefits; (2) The cap on child benefits affects only a subgroup of families with children. Therefore, we can use the non-affected but similar families with children as a control group, rather than childless couples; (3) We are able to isolate the income effect and shut down the fertility response by looking only at families where at least one parent has consulted a doctor regarding sterilisation prior to the reform; (4) The 2011 cap on child benefits was repealed after being in place for a single year, which allows us to assess long term effects of a temporary income shock that was perceived to be permanent; (5) Detailed Danish data allow us to control for individual fixed effects and analyse mechanisms, heterogeneity, and timing of policy responses.

In an unpublished note, Almlund (2018) discusses fertility effects of the 2011 reform in child benefits in Denmark. Almlund (2018) finds that the reduction in child benefits significantly decreases the probability of having a third or higher order child. Our focus is on the labour supply effects of the reform, but fertility is, of course, a crucial determinant of labour supply. Therefore, we separate the labour supply responses from the fertility response by controlling for family characteristics and by exploiting data on medical consultations regarding sterilisation, a common procedure in Denmark.

3. Child benefits in Denmark

In this section, we briefly outline the Danish child benefits system, the 2011 reform and its repeal, with further details reserved to Appendix A. Since 1986, families with dependent children have received child benefits from the government in Denmark. The level of child benefits paid to each family depends on the number of children as well as on the age of each child, with younger children assigned the highest benefit level. Child benefits are paid quarterly directly to the child's mother, or father if no mother is present. The first payment is made at the beginning of the quarter following a child's birth.² 2011 per-child benefits are listed in Table 1. For families with multiple children, the benefits per child are simply added together. For example, a family with three children of ages 1, 5, and 8 would receive $16,992 + 13,452 + 10,584$ DKK = 41,028 DKK (\approx 5,900 USD) in annual child benefits in 2011. Importantly, child benefits are not subject to income taxation in Denmark, where tax rates otherwise are relatively high. Fig. 1 shows the distributions of child benefit payments in 2011 (ignoring the influence of the cap). All those to the right of the red line in this figure were affected by the reform.

² All child benefit payments are made on dates 20 January, 20 April, 20 July, and 20 October every year.

Table 1

Yearly per-child benefit levels in 2011.

Age of child:	0–2	3–6	7–17
Annual payment (DKK)	16,992	13,452	10,584

Notes: Yearly amounts of child benefits per child in 2011. Child benefits are paid in four quarterly instalments directly to the child's mother, or father if no mother present. Amounts are adjusted yearly to account for inflation. See Appendix A for details.

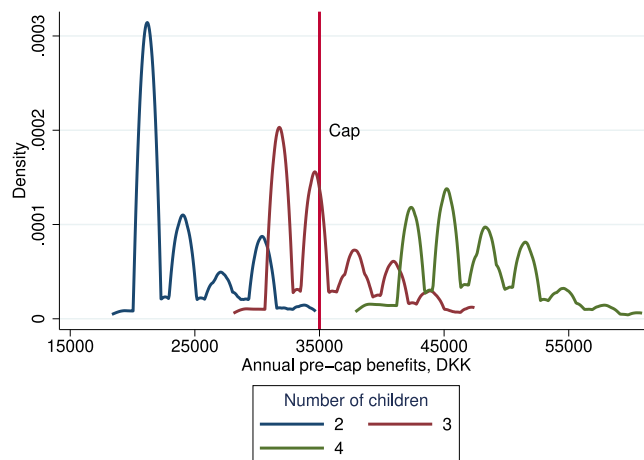


Fig. 1. Distributions of yearly child benefit payments.

Notes: Annual pre-cap benefits in 2011. Mean levels of benefits are: 23,669 DKK for 2-child families, 35,027 DKK for 3-child families, and 46,511 DKK for 4-child families. Cap at DKK 35,000 indicated. Child benefits are calculated using information on family composition from the population register BEF and verified using data on actual 4th quarter payments from BOBO/BOTL. Number of children refer to children eligible for child benefits, i.e. less than 18 years old. Although child benefits generally are paid in discrete amounts (see Table 1), in the quarter a child turns 18, child benefits are cut proportionally to the number of days under 18, which smooths the distribution. Observations below/above 7th/99th percentiles dropped due to confidentiality restrictions on Danish register data. Epanechnikov kernel density, bandwidth = 500.

3.1. Reform

In May 2010, in response to strong pressure on public finances after the global financial crisis, the Danish government announced the introduction of a cap, set at 35,000 DKK (2011-level, approx. 5000 USD) on total child benefits received by each family. The policy affected child benefit payments from January 2011 onwards. The policy is estimated to have affected 50,000 families. To smooth the income shock, child benefits were scheduled to be reduced evenly over 3 years: In 2011, child benefit payments would be reduced by one third of the amount exceeding 35,000, in 2012 by two thirds of the amount, and in 2013 by three thirds. Furthermore, the policy included a maximum reduction in benefits of approximately 12,000 DKK per year; a maximum that was to be gradually increased over the coming years (see details in Appendix A). These rules imply that the reduction in child benefits in 2011 was generally proportional to the total expected future loss of benefits for each family in the years to follow.

The 35,000 DKK-cap on child benefits did not directly affect one-child families. Although the reform did not directly affect the financial situation of two-child families, it changed the marginal child benefit received for having an additional child for two-child families and above. In other words, two-child families would receive a lower amount of child benefits for a potential third child after the reform. The cap did directly affect the unearned income of some three-child families with young children, specifically of families with at least one child under the age of 3 or at least two children under the age of 7. The reform changed the unearned income of all families with four or more children. Fig. 2 shows the total child benefits received by families of different structures from 2008 to 2016 after the introduction of the reform. Figure C.15 in

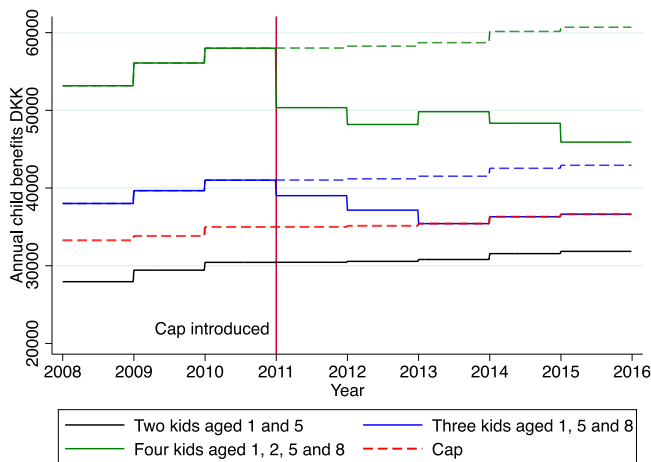


Fig. 2. Introduction of reform.

Notes: This figure illustrates child benefit policy as announced in 2010 and implemented from January 2011. DKK 35,000 cap introduced in 2011 in dashed red, family counterfactuals (benefits if no cap) are also dashed. Child benefits level are inflation-adjusted yearly. Notice the gradual (and non-linear) phase-in of the reform; it would be fully phased-in by 2020 (see details in Appendix A). Child age composition is fixed across years; we compare the effect of the reform on families with children of similar ages across years. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Appendix C shows the distribution of child benefits losses in 2011 for families with different numbers of children.

3.2. Repeal

At the end of November 2011, a newly elected left-wing government announced that from 2012, the cap on child benefits would be abolished due to its outsized impact on the income of large families. As a result, the loss in child benefits for those with large families was ultimately restricted to the year 2011, with 2012 payments returning approximately to 2010 levels for the affected families. Fig. 2 shows the reform as it was originally enacted in 2011, and hence expectations about future child benefit payments at the time of the introduction of the reform. In contrast, Fig. 3 illustrates both the introduction and the repeal of the reform, which was the policy schedule implemented after the election of the new government in November 2011. Considering the effect of the reform on various family compositions, we see that families with more children were relatively more affected by the reform. However, until November 2011 (when the repeal of the cap on benefits was announced), families faced a long-term reduction in income that would affect them until their children were old enough to no longer qualify for child benefits exceeding a total of 35,000 DKK.

See further details on the Danish child benefits system, and the 2011-reform and its repeal in Appendix A.

4. Hypotheses

In Appendix B, we outline a simple static model in which parents derive utility from having children, from consumption, and from leisure. Children are associated with a financial cost, which is in part mitigated by the child benefit system. Compared to in-work child benefits (e.g. the EITC in US), the analysis of the cap on child benefits in Denmark is more straightforward as there is no substitution effect between labour supply and child benefits. In our case, we can isolate the income effect. In addition to income from child benefits, parents can either spend time in the labour market, earning an hourly wage, or they can receive a fixed transfer from the government.

Our simple model has a number of key implications. With child benefits, the cost of an additional child is reduced, and hence, fertility

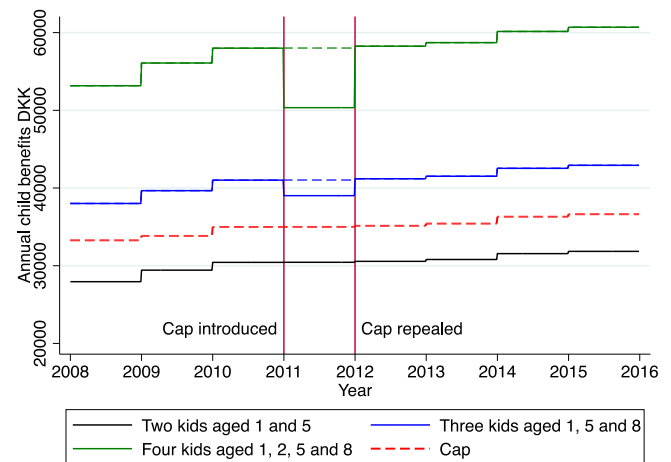


Fig. 3. Repeal of reform.

Notes: This figure illustrates child benefit policy as announced in November 2011 and implemented from January 2012. DKK 35,000 cap introduced in 2011 in dashed red, family counterfactuals (benefits if no cap) are also dashed. Child benefits level are inflation-adjusted yearly. See details on the child benefit reforms in Appendix A. Child age composition is fixed across years; we compare the effect of the reform on families with children of similar ages across years. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

is expected to increase with the level of child benefits. At the same time, an increase in unearned income generates income effects. The marginal utility of income, and therefore, of work decreases, predicting a lower level of labour supply when child benefits are implemented. As a result, in terms of labour supply, the introduction of a cap on child benefit is predicted to increase labour supply at the intensive margin for individuals in work and with large numbers of children. Secondly, we expect to see an increase in labour supply at the extensive margin for individuals not working and with large numbers of children. At the time of the reform, child benefits were exempt from taxation and determined independently of earnings and income. Thus, all labour supply effects are due to income effects alone.

In terms of fertility, the reform should not affect individuals with few children and those with a target household size that includes only a few children. In contrast, we do expect to see reduced fertility among individuals with a target household size that includes more children. Thus, our model reveals that families affected by the reform primarily can respond to the cap on child benefits through two channels: (1) increase their labour supply, and (2) decrease their target household size. A decreased target household size can again affect labour supply. In the sections that follow below, we will test the predictions of the model set out here. We are particularly interested in the labour supply effects of the reform on families that have already reached their target household size. With traditional data sources, determining whether or not families have reached their target household size is difficult, especially at the population level. However, in our case, we can proxy this by observing parents' medical consultations on sterilisations. We use such a consultation as an indicator that a household has already reached their target household size.

Although the cap on child benefits was intended to be permanent when it was announced and introduced, it was repealed after being in place for just a year as a new government was elected. However, despite the repeal of the cap, there are at least two reasons to expect permanent labour effects of the reform.

First, if we consider policy responses on the extensive margin, we may find that parents enter the labour market and carry the costs of entry when their non-market income decreases as child benefits are capped (Cogan, 1981). Costs upon entry into the labour market can include both search costs when finding a job and costs of searching for childcare. If parents also face a cost when exiting the labour market,

parents are likely to remain in the labour market after the repeal of the reform. Costs upon exit may include time spent on claiming unemployment benefits, social stigma, interacting with social workers, etc. We model this as a switching cost in the model outlined in Appendix B.

Secondly, the reform introduced uncertainty about the future levels of child benefits. Prior to the 2011 cap, child benefits in Denmark had not been subject to any substantial cuts since their introduction, only increases. With the 2011 cap, child benefits suddenly attracted (negative) political attention, and parents could likely be subject to another cut in benefits if the balance of power in the parliament shifted again.³ If parents are risk averse, uncertainty about future levels of child benefits is likely to affect labour supply at both the intensive and extensive margins in the long run for families that were subject to the 2011 cap.

A more complete theoretical framework should consider interactions between the two parents. However, for both our treatment and control groups, fathers' labour supply is already very high in the pre-treatment period. Therefore, we argue (and our results support) that fathers have few options to increase their labour supply further. In that sense, we can consider women secondary earners, and we argue that considering mothers' labour independently is sufficient in our context. The agent deciding the mother's labour supply could be considered to be the mother herself or the joint household; the intuition is similar, although a household bargaining model could further digest how couples reach a decision to increase maternal labour supply. However, in the context of this paper, we believe that the current model is sufficient to rationalise labour supply decisions.

5. Data, sample selection, treatment and control groups

5.1. Data and sample selection

To estimate the effect of the cap on child benefits, we rely entirely on administrative data from Denmark that are supplied by Statistics Denmark, covering the entire Danish population.

First, we construct a balanced panel with monthly observations of women and men covering the period September 2008 to December 2014. We exclude observations prior to September 2008 as a small change in the amounts of child benefits was implemented earlier in 2008. From Danish register data (BEF, OF), we can construct a monthly dataset with individual-level data on number of children, age of children, own age, gender, immigrant background, and parental leave status. Additionally, we have annual observations of marital status and cohabiting partners (BEF). Monthly hours worked are available from 2008 onwards in the BFL-register. Employers report monthly hours of work to the tax authorities together with monthly earnings. If an individual works more than one job in the same month, we aggregate hours worked from all jobs in the relevant month. This captures the total monthly hours worked.⁴

From the Danish administrative registers (BOBO/BOTI), we also have data on child benefit payments, but these registers only include data on payouts from the fourth quarter of each year from 2009 to 2014. However, from the population registers, we can construct precise

child benefits payments based on each family's number of children and the children's ages. Importantly, we can validate our measure of child benefits with the fourth quarter register data, enabling us to exclude families who receive irregular child benefits payments, e.g. because one or more children are in foster care.

Lastly, we obtain data on consultations with doctors regarding sterilisation procedures (from 1994 onwards from the registers LPRADM, SYSI, and SSSY). We construct monthly indicators that are equal to one if an individual has ever consulted a doctor regarding sterilisation. Sterilisation is a relatively common procedure to undergo in Denmark after finalising fertility (see Figure C.26 in Appendix C). One of the main advantages of the administrative data is that we observe personal identifiers of both parents and partners, and thus we can also construct family-level variables, which allows us to construct an indicator equal to one if either partner in a couple has consulted a doctor regarding sterilisation.

The resulting dataset is a balanced panel with monthly observations of individuals from September 2008 to December 2014 with data on their labour supply, demographics (including family composition), child benefits payments, and family sterilisation consultations.⁵

Next, we impose a number of sample selection criteria. First, we exclude people who are less than 25 years old by September 2008 or at least 60 years old by December 2014. This criterion ensures that our sample is in prime working age throughout our sample window. We also drop parents of children to whom child benefits are paid irregularly. Next, we exclude immigrants and children of immigrants because the eligibility to receive child benefits changed for immigrants in the sample period, and also because women's labour supply elasticities may depend on cultural background (see e.g. Fernández and Fogli, 2009).

For our main analysis, we also drop individuals who are self-employed at any point within our sample window because we do not observe hours worked for this group.⁶ We include self-employed in a robustness check considering earnings, rather than hours worked. For that exercise, we consider income measures from the annual income register (IND) as income from self-employment is not measured monthly.

5.2. Defining treatment and control groups

We group our families into those who were affected by the policy (treatment group) in the first two quarters of 2011, i.e. the families that experienced a strictly positive reduction in child benefits due to the policy, and those who were not affected by the policy (control group) in the first two quarters of 2011. We define the treatment/control groups based on the first two quarters of 2011 to avoid selection in and out of the treatment group. Child benefit payments in a given quarter are based solely on the number and ages of the children in a family in the *previous* quarter. As the child benefit reform was announced in late May 2010, children conceived before the announcement of the policy may have been born until March 2011, and thus will affect the child benefits paid in the second quarter of 2011. In other words, children born after March 2011 will almost definitely have been conceived after the announcement of the reform, and they will affect child benefits payments from the third quarter of 2011 onwards.

When choosing our treatment and control groups, we want to select groups that are as similar as possible in all ways except for their treatment under the policy. In general, the policy affected families with more children and with younger children. Families with 3 children under 18 are affected by the policy if they have at least one infant (0–2 years old) or two children in the age bracket 3–6 years. All families

³ A way to measure the level of public debate on child benefits is to look at the number of newspaper articles published on the subject. Numbers from the Danish media database Informedia show that an average of 108.8 articles per year mentioned child benefits (DA: "børnepenget") between 2007–2009 in national newspapers. From 2010–2014, the average number of articles per year increased to 280.2.

⁴ During parental leave in connection with childbirth, parents are entitled to full salary from their employers for part of the spell of parental leave. For some employees with fixed contractual hours of work, hours in BFL are not adjusted during spells of parental leave. As we observe days of leave, we can adjust the hours of work accordingly. See e.g. Ray et al. (2010) for a discussion of parental leave policies in Denmark.

⁵ Because we balance the panel, individuals who are only in the Danish population for part of the sample period are dropped, e.g. if an individual lives abroad for a year during the sample window.

⁶ 13.0% (29,144 individuals) of men and 7.3% (18,940 individuals) of women are dropped due to being primarily self-employed in any year between 2008 and 2014.

with 4 children under 18 are affected by the policy. To have young children in our control group after the introduction of the policy, we therefore cannot limit our sample to families with 3 children or more under 18. Instead, we limit our analysis to the individuals who in January 2011: (1) Had at least 2 children under 18.⁷ (2) Had no more than 4 children in total. (3) Were married or in a cohabiting couple. Therefore, parents with young children are included in both the control and treatments groups, and 4-child families are also present in the control group if at least one child is older than 18.

After we impose these additional selection criteria, we balance the panel, excluding people who are not present in the Danish population registers and in our sample for the entire sample period, September 2008 to December 2014. We are left with 28,968 treated women and 22,477 treated men. The difference in sample size between women and men is the result of the sample selection criteria: Compared to women, men are more likely to be self-employed, and men are typically older when they have children.

In Appendix C, we report summary statistics for the various samples in Tables C.8 to C.11, and we also show the distributions of hours worked and age in Figures C.11 to C.14. Generally, both the treated women and men are younger, have less labour market experience, are better educated, and are more likely to have twins or triplets when compared to the control groups. On average, the youngest child is also younger in the treatment group. This is expected, given the design of the cap. The 35,000 DKK-cap on child benefits changed the unearned income for three-child families with young children, specifically for families with at least one child under the age of 3 or at least two children under the age of 7 years old. The reform changed the unearned income for all families with four or more children. In other words, three-child families are included in the control group if their children are slightly older, and in the treatment group if they are slightly younger. This will translate into differences in the average age of youngest child and in the average age of the parents. The fact that the treatment group women on average are 3 years younger than women in the treatment group, but have 4.1 fewer years of labour market experience is a result of women in the treatment group having more children, and therefore they have spent more time on parental leave. In our empirical strategy, which we outline in the following section, we specifically address these differences in characteristics between our treatment and control groups, and next, we compare pre-trends in outcomes for the different groups.

Lastly, note that “Family’s yearly benefits lost in 2011” are not exactly equal to zero for the control group. This is because we define our treatment group as those who were affected by the policy in the first two quarters of 2011 to avoid selection into the treatment group; this is further detailed above.

6. Empirical strategy

Our aim is to test whether the temporary introduction of the child benefit cap affected labour supply, with a particular focus on the direct income effect rather than any indirect effects through fertility. Our first set of specifications will simply group families into those who saw their child benefits reduced due to the reform (treatment group) and those who were not affected by the reform (control group) as described above. By classifying families in this way, we can estimate the effect of the reform by using a binary treatment indicator. This is a simplification of the reform, but it is one that is particularly useful for a graphical analysis and provides first evidence of the direction of a potential effect. We will then extend our analysis and consider the treatment as continuous. For all of the analyses below, we run the analyses separately for women and men and compare the results.

⁷ We also require that both children were registered in the Danish population register BEF in 2011 (or 2012 for children born in 2011) with the relevant parent ID.

6.1. Binary treatment

We undertake two sets of analyses using our binary treatment indicator. First, we analyse treatment dynamics using a fully generalised difference-in-differences setup to determine whether or not parallel trends are a reasonable assumption for our further analyses. Second, we apply a standard difference-in-differences model, but considering three time periods, namely: (1) before introduction of reform, (2) after introduction of reform, and (3) after repeal of reform.

6.1.1. Generalised difference-in-differences

As parallel trends for the treatment and controls groups is the main identifying assumption for our further analyses, we start by estimating the following generalised difference-in-differences model:

$$Y_{it} = \beta_0 + \sum_{j=0}^T \beta_1^j 1[j = t] + \sum_{j=0}^T \beta_2^j 1[j = t] * W_i + \beta_3 X_{it} + m_i + \alpha_i + \epsilon_{it} \quad (1)$$

Where t is time in year-quarters, Y_{it} is an outcome of interest for person i in time t (e.g. hours worked or employment), W_i indicates whether i is in the treatment group (affected by benefit cut in Q1 and/or Q2 2011), and X_{it} is a set of time-varying family controls. $\sum_{j=0}^T \beta_1^j 1[j = t]$ is a set of year-quarter fixed effects, and m_i are 8 calendar month fixed effects that account for within-quarter seasonality (they are not year*month fixed effects; last calendar month of each quarter omitted to avoid collinearity with year-quarter fixed effects). Lastly, α_i are individual fixed effects. Note that a treatment group indicator without a time interaction is not included, as treatment group membership is fixed over time, and thus, absorbed by the individual fixed effects. We cluster standard errors at the individual level.

β_2^j are the parameters of interest. They yield the estimates of the differences in time trends between the treatment and control groups and will allow us to inspect both potential differences in pre-trends and the dynamic effects of the reform on outcome Y_{it} .⁸

6.1.2. Difference-in-differences

After inspecting pre-trends both in the actual levels of hours worked of our treatment and control group as well as in our generalised difference-in-differences results from the specification above, we conclude that the parallel trends assumption is full-filled — see the discussion of our results below. We also estimate our next specification with a binary treatment indicator for whether or not person i was affected by the cap on child benefits. Our main specification is:

$$Y_{it} = \beta_0 + \sum_{j=0}^T \beta_1^j 1[j = t] + \beta_2 W_i * Intro_t + \beta_3 W_i * Repeal_t + \beta_4 X_{it} + m_i + \alpha_i + \epsilon_{it} \quad (2)$$

The notation is analogous to that of Eq. (1), except that $Intro_t$ is an indicator equal to one in all periods after the policy introduction, and $Repeal_t$ is an indicator equal to one in all periods after the policy repeal.⁹

$\beta_2 > 0$ yields the average treatment effect from 2011, and $\beta_3 > 0$ the average additional treatment effect from 2012 to 2014. However, as will be discussed in Section 7, the treatment effect gradually increases throughout 2011 and remains constant after the repeal in 2012.

⁸ In addition to visual inspection of pre-trends, for our preferred specification (model C, described in Section 6.3), we test pre-trends similar to Meyer (2010). For the coefficients on the pre-period quarter indicators, we find that we cannot reject both the hypothesis that they are jointly equal to zero, and the hypothesis that they are jointly equal to each other at the 5%-level.

⁹ Our results are robust to controlling for common time trends with the non-interacted indicators $Intro_t$ and $Repeal_t$, instead of the dynamic controls $\sum_{j=0}^T \beta_1^j 1[j = t]$. We focus on the specification in Eq. (2) to be consistent with that of our generalised difference-in-differences approach in Eq. (1).

Therefore, $\beta_2 > 0$ will underestimate the full treatment effect as the coefficient will be an average over the four quarters of 2011, including the small effect in Q1 2011 and the full reform effect in Q4 2011. Because the treatment effect remains constant after the repeal of the reform, we focus on the following specification where we drop the interaction term between the treatment group and the repeal of the reform:

$$Y_{it} = \beta_0 + \sum_{j=0}^T \beta_1^j 1[j = t] + \beta_2 W_i * Intro_t + \beta_4 X_{it} + m_i + \alpha_i + \epsilon_{it} \quad (3)$$

Keep in mind that $Intro_t$ is an indicator equal to one in all periods after the policy introduction. Therefore, our main hypothesis is that $\beta_2 > 0$. Again, we cluster standard errors at the individual level. We report the results from estimating Eq. (3) in the main text of the paper, and reserve the estimates from Eq. (2) to Appendix C.

6.2. Continuous treatment

In practice, the magnitude of the policy treatment is continuous. Therefore we also estimate the following specification including a continuous treatment measure:

$$Y_{it} = \beta_0 + \sum_{j=0}^T \beta_1^j 1[j = t] + \beta_2 V_i * Intro_t + \beta_4 X_{it} + m_i + \alpha_i + \epsilon_{it} \quad (4)$$

Again, the notation is analogous to that of Eq. (1), but V_i is the reduction in child benefits in thousand DKK in 2011 due to our reform for i 's family if i is in the treatment group (affected by benefit cut in Q1 and/or Q2 2011) and zero otherwise. Again, we cluster standard errors at the individual level. β_2 is the parameter of interest, which yields an estimate of the change in Y_{it} for every 1000 DKK lost in child benefits due to the reform.

We do not consider the actual level of child benefits, but only the reduction in child benefits due to the reform. Individual fixed effects and family controls capture the factors that otherwise determine the total level of child benefits. Also for the continuous treatment, we estimate a version with an indicator for the repeal of the reform similar to Eq. (2); these estimates are reported in Appendix C.

6.3. Identification: Fertility and treatment effects

Due to the fertility effects of the reform described in Almlund (2018), we can potentially attribute a significant share of a treatment effect of the reform to decreased fertility. To disentangle the effects of changed fertility on labour supply from the income effect, we run three versions of the analyses described above, where we add different control variables and additional sample selection criteria in the following three models:

(A) Parental controls:

Includes only individual FEs and parental age FEs. Parental age FEs include each quarter of age in order to control non-parametrically for age effects of parents.

(B) Parental controls + Family controls:

Same as model (A), but also includes age of youngest child FEs and total number of children FEs.¹⁰ Age of youngest child FEs include each quarter of age in order to control non-parametrically for age effects children.

(C) Parental controls + Family controls + Sterilisation restriction + No young restriction:

Same as model (B), but limited sample including only families in which at least one parent has consulted a doctor regarding sterilisation prior to the announcement of the policy. Also excluding families with young children (below the age of 3) throughout the sample period.

In model (A), we estimate the combined effects of decreased fertility and the income effect of losing child benefits due to the reform. However, the results from model (A) should be interpreted with caution and as an upper bound. When comparing those who were affected by the policy in the first two quarters of 2011 with those who were not, the treatment and control groups are fixed over time. An issue with this approach is that the probability of having reached their target household size can vary between households in the treatment and control groups. As treatment assignment is partly determined by having young children (3-child families), having reached their target household size will result in fewer parents being on parental leave, and thus, they increase their labour supply. In that sense, part of the estimated effect from model (A) will be mechanical; this is confirmed by our placebo test in Appendix C.7 where we apply our empirical strategy at a point in time where families were not affected by the cap on child benefits.

We are primarily interested in estimating the labour supply effect of the reform net of the fertility response and focus on the income effect of the policy shift. We take two approaches to identify this effect. First, in model (B), we control for a changing composition of children over time. Separately controlling for the number of children and age of youngest child specifically addresses the threats to identification in model (A), namely a difference in the probability of having reached target household size and different fractions of parents on parental leave at the time of introduction of the reform. However, our estimates in model (B) may be downward biased as our control group may respond to the cap on child benefits by delaying or changing fertility decisions, even though their incomes were not directly affected by the reform at the time of introduction. For example, a family with two children may respond to the policy by not having a third or fourth child as they realise that they will receive a lower level of child benefits for the additional children. Our placebo test in Appendix C.7 also suggests that we should address this source of bias. In addition, the family controls in model (B) may also be affected by the treatment through a potential fertility effect, which could lead to bias in our inference through a 'bad control' problem (Angrist and Pischke, 2008).

Rather than pursuing an IV-strategy to eliminate these potential sources of bias, we first control for fertility-related family characteristics. Secondly, we exploit data on sterilisation, focusing on families that have finalised fertility prior to the announcement of the reform in model (C).¹¹ In model (C), we shut down the fertility response completely by only considering families where one or both parents have consulted a doctor regarding sterilisation prior to the announcement of the reform and who do not have young children.¹² Model (C) therefore provides our best estimates of the income effect of the reform on labour supply.

¹¹ Existing papers have estimated the effects of child/family benefits using various empirical strategies. Many papers use a difference-in-differences strategy similar to ours, but comparing families with and without children. Importantly, Dahl and Lochner (2012) have suggested an IV-approach when they exploit changes in the EITC over time to estimate the impact of family income on children's math and reading achievement. The IV-strategy primarily aims to alleviate omitted variable bias and bias from measurement error affecting income, something that is not a concern in our setting where levels of child benefits are determined independently of income.

¹² If we impose the sterilisation consultation sample restriction alone, we will have families with young children in the pre-reform period (before parents consult a doctor regarding sterilisation), but no young children in the post-reform period (after parents consult a doctor regarding sterilisation). Thus, we would be comparing very different types of families if we do not also limit our sample to families without very young children. Families with children aged 3 and above are observed throughout the period of the panel, also when conditioning on sterilisation consultation prior to the announcement of the policy. Hence, it is necessary to impose the two criteria together for correct identification of the reform effect net of a fertility response.

¹⁰ Of children registered in the Danish population register BEF in each year from 2008–2014.

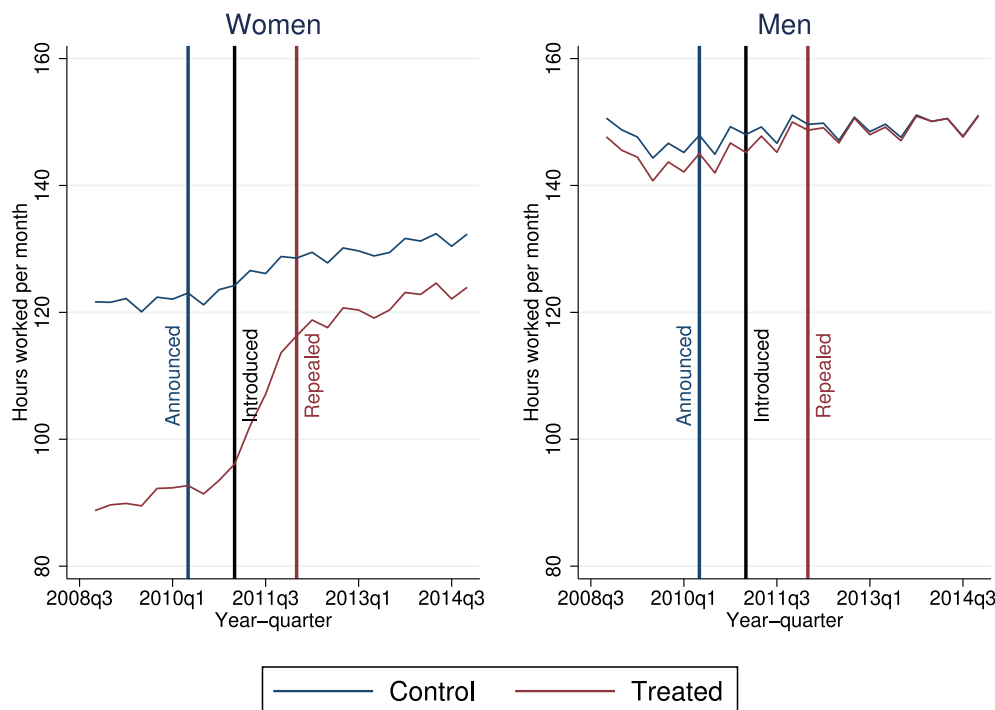


Fig. 4. Hours worked per month for women and men.

Notes: Average hours worked per month for women and men respectively, including individuals working zero hours. Quarter averages calculated from monthly observations from October 2008 to December 2014. Monthly hours worked are available from the BFL-register. See details on the estimation sample in Tables C.8 and C.9.

The intention of focusing on families where at least one parent has consulted a doctor regarding sterilisation is to make preferences for additional children comparable across the treatment and control groups, which otherwise are likely to differ between families with different numbers of children. However, even if preferences for additional children are fixed and comparable across the two groups, preferences for leisure/time with children may also differ between families with different numbers of children. Therefore, we also include families with 4 children in the control group if at least one, but no more than 2 of their children are over 18. Preferences for leisure/time with children may vary over time, and particularly with the age of youngest child, and therefore, we control for age of youngest child in models (B) and (C). If the treatment group, conditional on these controls, still has stronger preferences for leisure/time with children than the control group due to the higher number of children, our estimated labour supply responses should provide a lower bound for the group of families with fewer children.

7. Results

7.1. Descriptives

Firstly, we inspect the raw trends in hours worked per month for our treatment and control groups. In Fig. 4, we see that the trends in hours worked are very similar for the two groups before the introduction of the reform as well as after its repeal. This serves as a first indication of the parallel trends assumption being full-filled for our treatment and control groups. However, we see a large increase in hours worked per month for our treatment group during the reform period, roughly 20 h per month. Notice that the response to the reform is gradual after its introduction; mothers do not respond immediately after the announcement of the reform. This is in line with Hotz and Scholz (2006, p. 42) who find that “the employment responses to EITC policy changes occur with a lag of one or two years”.

We also see a very small increase in hours for our control group after the introduction of the reform. Although the control group is not directly affected by the cap on child benefits, families with two children can be indirectly affected if they were planning to have a third child before the introduction of the reform. If they have a third child, their child benefits for this child will be reduced after the introduction of the reform. Therefore, families may delay or reconsider having a third child due to the reform. Therefore, if fertility is fixed, we should not observe this effect. We use our sterilisation subsample to confirm this, see Figure C.17 in Appendix C.

Secondly, in Fig. 5, Panel (a), we see a general decrease in fertility rates. In Fig. 5, Panel (b), we consider number of births per 31-year-old-woman (median age of mother at birth), which suggests that the general declining trend in fertility rates partly is due to smaller cohorts of potential mothers (based on age). We also see that the number of births decreases disproportionately in families with two or more children after the introduction of the policy. This is in line with the results of Almlund (2018). However, the decline in high-order births appears to have started already from 2008 to 2009, at the onset of the Great Recession. We keep this in mind throughout our analysis by estimating the effects of the reform both with and without controls for fertility-related characteristics. The main focus of our analyses is to estimate the effect of the cap on child benefits net of a potential fertility response.

7.2. Generalised difference-in-differences

After concluding that the raw trends in hours worked per month are very similar for our treatment and control groups before the policy announcement, we further inspect pre- and post-trends using our generalised difference-in-differences model. In Fig. 6, we show the estimated coefficients from Eq. (1) with the set of controls from model (A), where we only include individual fixed effects and parent age fixed effects. Again, we observe parallel trends for our treatment and control groups. Furthermore, for women, we find a large increase in labour supply for

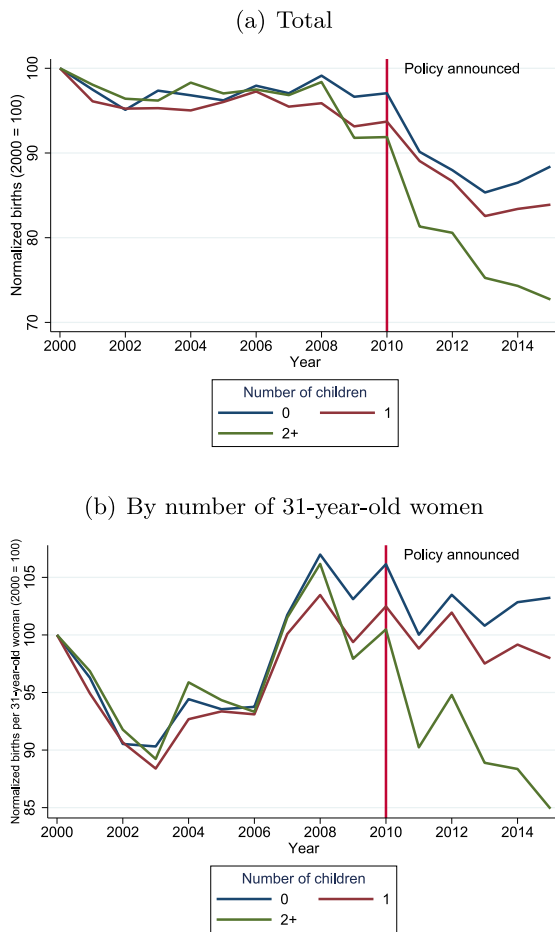


Fig. 5. Total number of births by pre-birth number of children.

Notes: For this exercise, we consider the number of births for the entire Danish population and not just the estimation sample from which 1-child families are excluded. Panel (b) considers number of births per woman turning 31 in the given year to address concerns about changing cohort sizes of mothers. Number of births refers to the number of newborns observed in the population register BEF. We normalise the number of births in 2000 to 100.

our treatment group relative to our control group, but only a small effect on men. Similar to the raw trends in hours worked per month reported above, we see that the effect increases through the reform period and remains stable after the repeal. The effects shown in Fig. 6 can be ascribed to both reduced fertility as well as a general income effect caused by the reduction in child benefits for our treatment group. However, as pointed out in Section 6.3, part of the estimated effect from model (A) is mechanical as preferences for additional children are likely to vary between the treatment and control groups. Therefore, the estimated effects should be interpreted with caution and as an upper bound. We address these concerns in the analyses that follow.

In order to separate the effect of reduced fertility from the general income effect of the reform – and to address the issue of different preferences for additional children across the treatment and control groups – we control for fertility-related family characteristics in Fig. 7. We see a smaller, but still significant effect on women's labour supply both during the reform period and after its repeal. For men, we do not observe any reform effect.

In order to fully mitigate the fertility response to the policy – and to fix preference for additional children across the treatment and control groups – we now limit the sample to families where at least one parent had consulted a doctor regarding sterilisation prior to the announcement of the policy. Furthermore, we exclude families with young children throughout the sample period and unbalance the panel

in this sub-analysis. Compared to Fig. 7, we see slightly larger effect sizes for women in this subgroup of families in Fig. 8, but with the smaller subsample the coefficients are also less precisely estimated.

7.3. Difference-in-differences

After inspecting pre-trends and the dynamics in the treatment effect of the reform, we now move on to report our difference-in-differences estimates, i.e. the estimated parameters of interest from Eqs. (3) and (4). In Table 2, Panel A, the results using our binary treatment indicator mirror the previous figures. For women, we see a treatment effect of 15 h worked per month when considering model (A). The results from our placebo test in Table C.15 confirm that part of this effect is mechanical due to different preferences for additional children. After controlling for fertility-related family characteristics, we find an effect of 1.94 h worked per month. The estimates from the subgroup of families who have consulted a doctor regarding sterilisation and without young children confirm these positive effects on labour supply, yielding an estimate of the treatment effect of 2.69 h worked per month. For men, we do not find any treatment effect when appropriately controlling for fertility-related characteristics.

In Table 2, Panel B, we replace our binary treatment indicator with a continuous measure of child benefits lost due the reform in DKK divided by 1000. 1000 DKK corresponds to approximately 160 USD. Consider that an average family in our treatment group experiences a reduction in child benefits of approx. 2300 DKK (see Appendix C), but within our treatment group there is substantial variation in the amount of benefits lost, e.g. four-child families are particularly affected by the reform. Child benefits are tax free, and with the relatively high Danish taxes on labour market income, $\approx 50\%$, this corresponds to about 4600 DKK in labour market income. We find large and significant effects on the treated women. For every 1000 DKK lost due to the reform, the estimates from model (A) in Column 1, show that women work 4.894 h more per month on average. Excluding the mechanical effect and any fertility effects, we find an effect between 0.300 and 1.031 h per month for every 1000 DKK lost. With a tax rate around 50% and hourly wages of 200 DKK, the estimated effect indicates that women offset the income loss due to the reform roughly 1-to-1 when excluding the fertility response by earning an additional $(12 \times 1.031 \times 200) \times 0.5 = 1237.2$ DKK for 1000 DKK lost in child benefits. For men, we do not find significant effects of the reform on labour supply when controlling for fertility-related characteristics.

8. Extensive and intensive margins

In this section, we further analyse the increase in average hours worked for women in the treatment group. In particular, we want to separate adjustments along the intensive and extensive margins of labour supply. As a first step, in Fig. 9, we map out the average hours of work per month for employed people only, i.e. we exclude people working zero hours.

From Fig. 9, it appears that there is only a small break in the trend of average hours worked per month for women in the treatment group after the reduction in child benefits. For men, the gap in hours worked between the treatment and control groups closes during the treatment period. Thus, we have some indication of men's labour supply being affected at the intensive margin during the treatment period.

In Fig. 10, we map out the participation rates for women and men, respectively. We see a jump in the participation rate for the treated women, but not for men. From Figures C.18 and C.19 in Appendix C, we see that the jump in the participation rate for treated women appears in both part-time and full-time jobs. We also see a small increase in the participation rate of women in our control group after the introduction of the reform. Although the control group is not directly affected by the cap on child benefits, families with two children would be affected if they had a third child. This may delay or change the decision to have

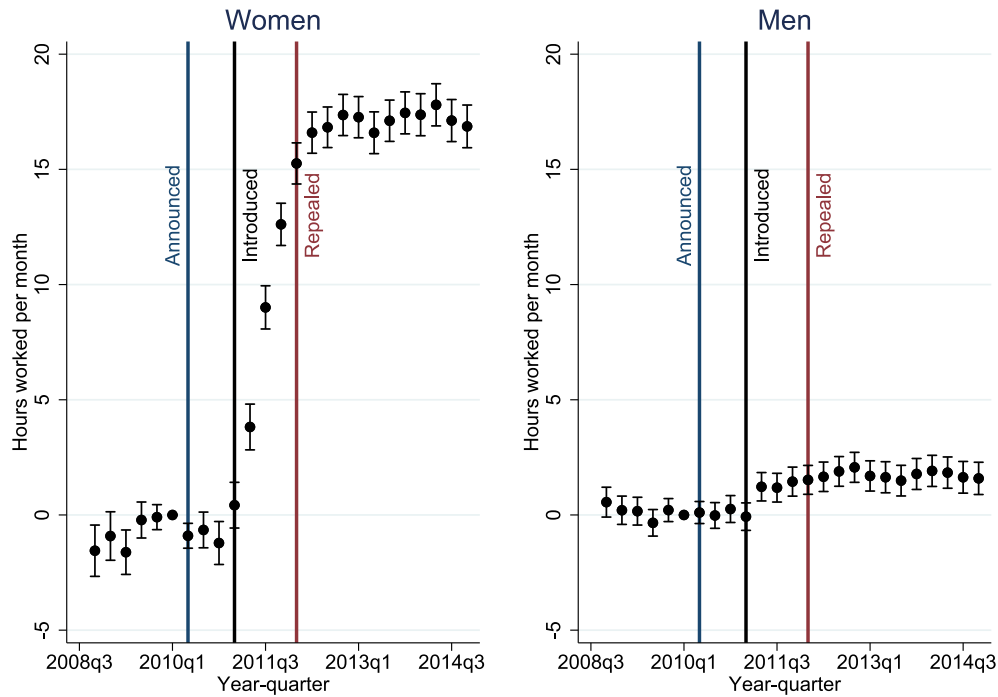


Fig. 6. G-DiD: Hours worked per month. Parental Controls.

Notes: Dependent variable (y-axis): hours worked per month. This figure shows the estimated coefficients from Eq. (1), model (A), as described in Section 6.3. Estimated on monthly observations from October 2008 to December 2014. Monthly hours worked are available from the BFL-register. See details on the estimation sample in Tables C.8 and C.9. Parental controls include: individual FEs and parental age FEs. 95%-confidence intervals indicated.

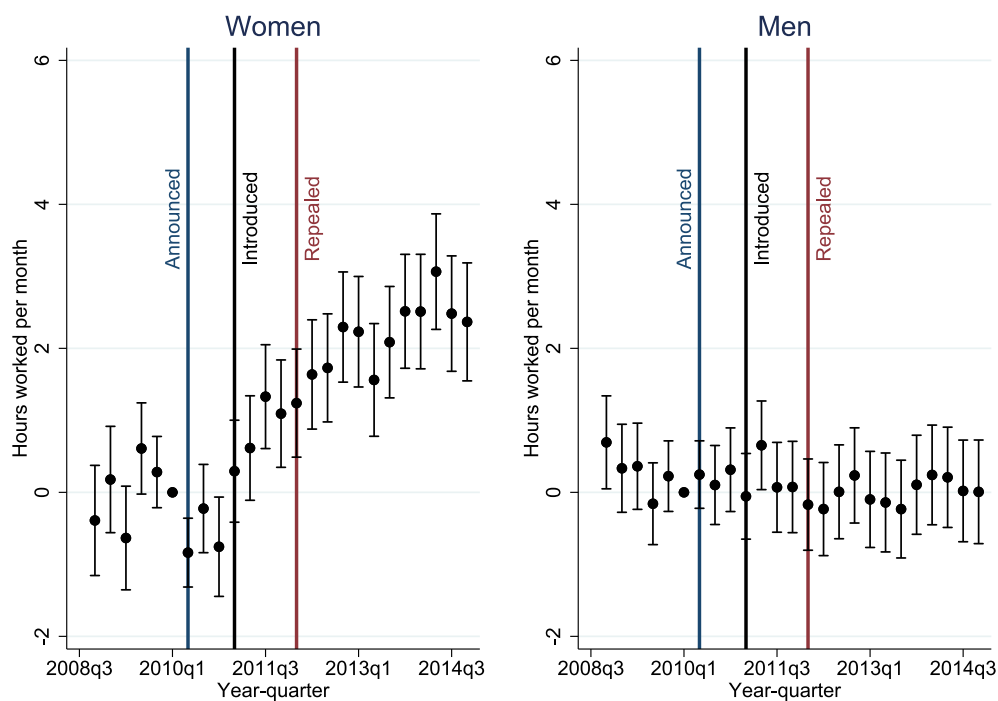


Fig. 7. G-DiD: Hours worked per month. Parental Controls + Family Controls.

Notes: Dependent variable (y-axis): hours worked per month. This figure shows the estimated coefficients from Eq. (1), model (B), as described in Section 6.3. Estimated on monthly observations from October 2008 to December 2014. Monthly hours worked are available from the BFL-register. See details on the estimation sample in Tables C.8 and C.9. Parental controls include: individual FEs and parental age FEs. Family controls include: age of youngest child FEs and number of children FEs. 95%-confidence intervals indicated.

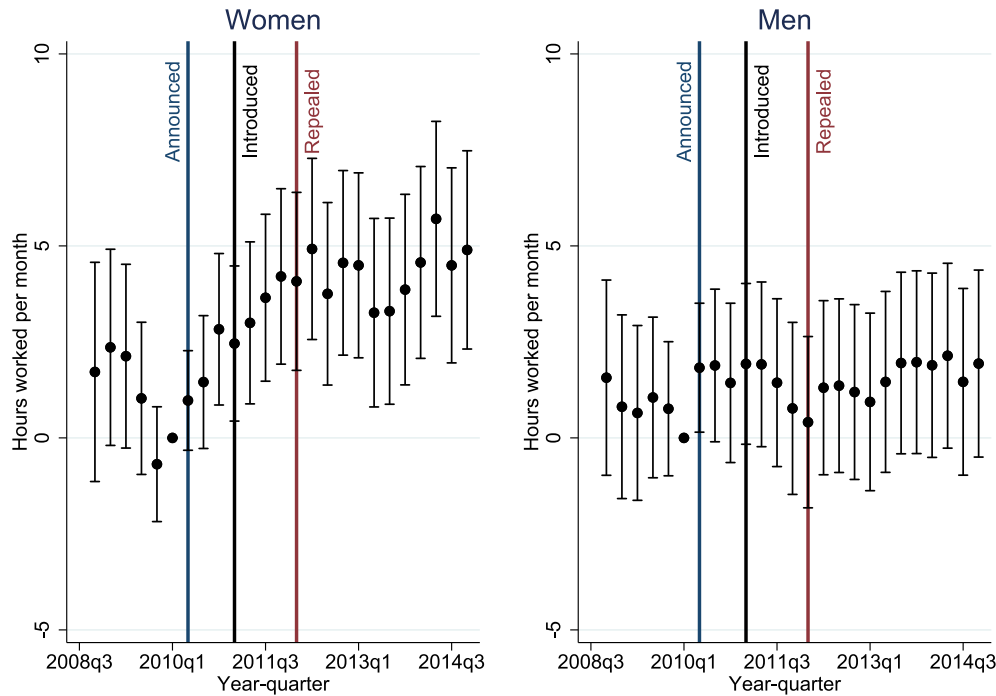


Fig. 8. G-DiD: Hours worked per month. Parental Controls + Family Controls + Sterilisation restriction + No young restriction.

Notes: Dependent variable (y-axis): hours worked per month. This figure shows the estimated coefficients from Eq. (1), model (C), as described in Section 6.3. Estimated on monthly observations from October 2008 to December 2014. Monthly hours worked are available from the BFL-register. See details on the estimation sample in Tables C.10 and C.11. Parental controls include: individual FEs and parental age FEs. Family controls include: age of youngest child FEs and number of children FEs. Sterilisation restriction + No young restriction: only including families where at least one parent has consulted a doctor regarding sterilisation prior to the announcement of the policy, and excluding families with young children (below the age of 3 throughout the sample period). 95%-confidence intervals indicated.

Table 2
Difference-in-differences: Hours of work per month.

Dependent variable: Hours worked per month						
	Women			Men		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Binary treatment						
1.Intro × 1.Treat	15.01*** [0.277]	1.944*** [0.243]	2.688** [0.846]	1.398*** [0.186]	−0.182 [0.196]	0.307 [0.685]
R ²	0.517	0.619	0.663	0.460	0.465	0.514
Panel B: Continuous treatment (reduction in 1000 DKK)						
1.Intro × c.Treat	4.894*** [0.101]	0.300*** [0.0897]	1.031** [0.332]	0.532*** [0.0672]	−0.0158 [0.0713]	0.0758 [0.248]
R ²	0.517	0.619	0.663	0.460	0.465	0.514
N	18 153 075	18 153 075	3 205 408	14 655 450	14 655 450	2575379
Female	1	1	1	0	0	0
Parental controls	1	1	1	1	1	1
Family controls	0	1	1	0	1	1
Sterilisation	0	0	1	0	0	1
Ex. young	0	0	1	0	0	1

Notes: Dependent variable: hours worked per month. This table reports the estimated coefficients from Eq. (3) in Panel A, and those estimated from Eq. (4) in Panel B. Estimated on monthly observations from October 2008 to December 2014. Monthly hours worked are available from the BFL-register. See details on the full estimation sample in Tables C.8 and C.9, and details on the sterilisation subsample with no young children in Tables C.10 and C.11. Parental controls include: individual FEs and parental age FEs. Family controls include: age of youngest child FEs and number of children FEs. Sterilisation restriction + Ex. young restriction: only including families where at least one parent has consulted a doctor regarding sterilisation prior to the announcement of the policy, and excluding families with young children (below the age of 3 throughout the sample period). The sets of controls are discussed in Section 6.3. Standard errors in brackets, clustered at the individual level, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

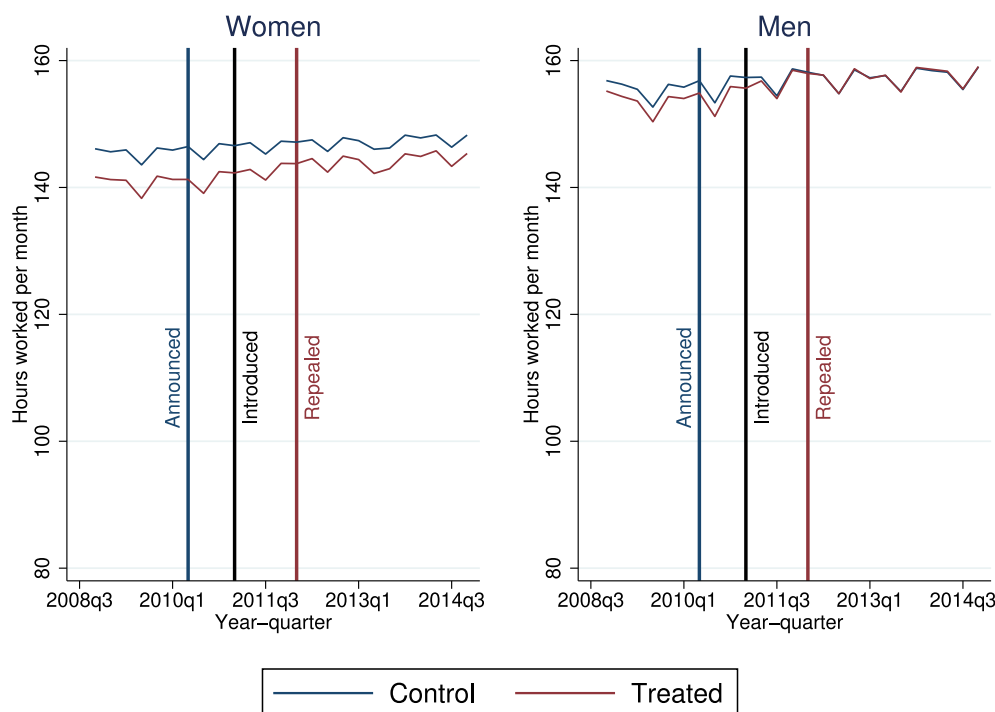


Fig. 9. Hours worked per month, excluding people working zero hours.

Notes: Average hours worked per month for women and men respectively, excluding individuals working zero hours. This illustrates changes in labour supply at the intensive margin. Quarter averages calculated from monthly observations from October 2008 to December 2014. Monthly hours worked are available from the BFL-register. See details on the estimation sample in Tables C.8 and C.9.

a third child. Again, we use our sterilisation subsample to address this concern.

These descriptive results suggest that women's response to the reform in child benefits happens through both the intensive and extensive margins of labour supply, whereas men may respond at the intensive margin. In the following subsections, we explore the responses at the extensive and intensive margins, respectively, in a casual setting using our difference-in-differences approach.

8.1. Extensive margin

We report the graphical generalised difference-in-differences results in Appendix C, see Figures C.20 to C.22. The graphical analyses again confirm parallel pre-trends. In Table 3, Panel A, the effects of the reform are reported using our binary treatment indicator. The results confirm that men do not respond to the reduction in child benefits at the extensive margin. For women, however, we observe a large response to the reform at the extensive margin. Including the mechanical response due to different preferences for additional children and any fertility response, we find an effect of 10.1 percentage points (Panel A, Column 1), but again, this estimate should be interpreted with caution and as an upper bound. When controlling for fertility-related characteristics, we estimate an effect ranging between 1.18 (Panel A, Column 2), and 1.36 percentage points (Panel A, Column 3) when also limiting our sample to those families in which at least one parent has consulted a doctor regarding sterilisation and who have no very young children, this corresponds to a 1.62% increase in labour force participation (0.0136/0.837, see Table C.11).

In Table 3, Panel B, the effects of the reform are reported using our continuous treatment measure. Again, these results confirm that men do not respond to the reduction in child benefits at the extensive margin. And again, we observe that women respond at this margin. When controlling for fertility-related characteristics, we estimate an effect ranging between 0.258 (Panel B, Column 2) and 0.489 percentage points (Panel B, Column 3) increased participation for every 1000 DKK lost due to the cap on child benefits.

8.2. Intensive margin

In order to assess labour supply responses at the intensive margin, we unbalance the panel and delete all monthly observations with zero hours of work. Again, we report the graphical generalised difference-in-differences results in Appendix C, see Figures C.23 to C.25. The pre-trends for women are less convincing for this sub-analysis, due to the changing composition of the sample as participation increases through the sample period, but keep in mind the parallel trends in Fig. 9. For men, the pre-trends are also parallel in Fig. 9, as well as in the generalised difference-in-differences setup.

Using our binary treatment indicator, the effects of the reform on the number of hours worked for people in employment are reported in Table 4, Panel A. Women also respond at the intensive margin by increasing their number of hours worked. The estimates are consistent across all specifications, although less statistically significant in Column 3. Controlling for fertility-related characteristics, we also find positive coefficients for men and with a magnitude that is economically meaningful, but they are statistically insignificant.

In Table 4, Panel B, we again report the effects of the reform on the number of hours worked for people in employment, but this time using our continuous treatment measure. These results confirm that both women respond to the reduction in child benefits at the intensive margin by increasing their number of hours worked. Notice that the estimates again are consistent across all specifications. Using our continuous treatment measure and controlling for fertility-related characteristics, we find an economically meaningful effect on men, but the effects either insignificant or only marginally significant (Panel B, Column 5).

9. Heterogeneity and earnings

Our main results show that women respond to the cap on child benefits by increasing their labour supply at both the extensive and

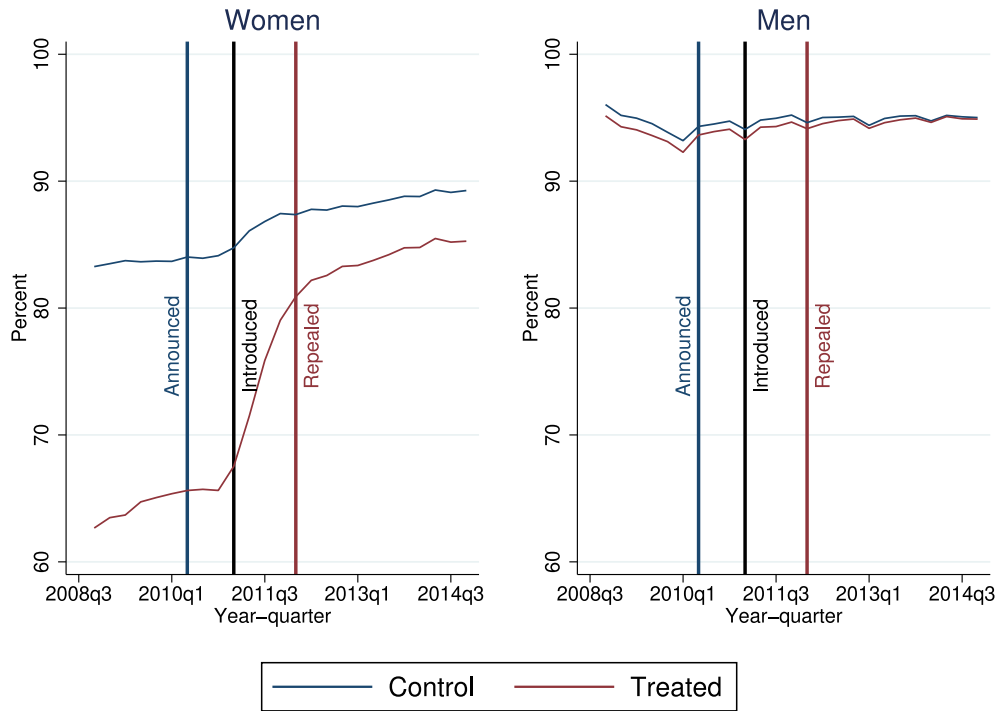


Fig. 10. Percent of women and men currently employed (i.e. working non-zero hours).

Notes: Percent of women and men currently employed (i.e. working non-zero hours). This illustrates changes in labour supply at the extensive margin. Quarter averages calculated from monthly observations from October 2008 to December 2014. Monthly hours worked are available from the BFL-register. See details on the estimation sample in Tables C.8 and C.9.

Table 3

Difference-in-differences: Employment indicator.

	Dependent variable: Employment indicator					
	Women			Men		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Binary treatment						
1.Intro × 1.Treat	0.101*** [0.00174]	0.0118*** [0.00149]	0.0136** [0.00513]	0.00120 [0.000979]	−0.00258* [0.00104]	0.000186 [0.00377]
R^2	0.466	0.602	0.638	0.467	0.468	0.525
Panel B: Continuous treatment (reduction in 1000 DKK)						
1.Intro × c.Treat	0.0335*** [0.000642]	0.00258*** [0.000558]	0.00489* [0.00203]	0.000480 [0.000353]	−0.000697 [0.000380]	0.000288 [0.00133]
R^2	0.466	0.602	0.638	0.467	0.468	0.525
N	18 153 075	18 153 075	3 205 408	14 655 450	14 655 450	2575379
Female	1	1	1	0	0	0
Parental controls	1	1	1	1	1	1
Family controls	0	1	1	0	1	1
Sterilisation	0	0	1	0	0	1
Ex. young	0	0	1	0	0	1

Notes: Dependent variable: Employment indicator (i.e. indicator for working non-zero hours). This table reports the estimated coefficients from Eq. (3) in Panel A, and those estimated from Eq. (4) in Panel B. Estimated on monthly observations from October 2008 to December 2014. Monthly hours worked are available from the BFL-register. See details on the full estimation sample in Tables C.8 and C.9, and details on the sterilisation subsample with no young children in Tables C.10 and C.11. Parental controls include: individual FEs and parental age FEs. Family controls include: age of youngest child FEs and number of children FEs. Sterilisation restriction + Ex. young restriction: only including families where at least one parent has consulted a doctor regarding sterilisation prior to the announcement of the policy, and excluding families with young children (below the age of 3 throughout the sample period). The sets of controls are discussed in Section 6.3. Standard errors in brackets, clustered at the individual level, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 4

Difference-in-differences: Hours of work per month, ex. zero hours.

	Dependent variable: Hours worked per month, ex. zero hours					
	Women			Men		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Binary treatment						
1.Intro \times 1.Treat	1.592*** [0.141]	1.319*** [0.154]	1.167* [0.531]	1.265*** [0.112]	0.171 [0.116]	0.394 [0.389]
R^2	0.476	0.480	0.545	0.333	0.343	0.382
Panel B: Continuous treatment (reduction in 1000 DKK)						
1.Intro \times c.Treat	0.506*** [0.0542]	0.350*** [0.0611]	0.556** [0.214]	0.490*** [0.0414]	0.0880* [0.0431]	0.0700 [0.147]
R^2	0.476	0.480	0.545	0.333	0.343	0.382
N	15 440 331	15 440 331	2 839 834	13 883 537	13 883 537	2441788
Female	1	1	1	0	0	0
Parental controls	1	1	1	1	1	1
Family controls	0	1	1	0	1	1
Sterilisation	0	0	1	0	0	1
Ex. young	0	0	1	0	0	1

Notes: Dependent variable: Hours worked per month, excluding those working zero hours. This table reports the estimated coefficients from Eq. (3) in Panel A, and those estimated from Eq. (4) in Panel B. Estimated on monthly observations from October 2008 to December 2014. Monthly hours worked are available from the BFL-register. See details on the full estimation sample in Tables C.8 and C.9, and details on the sterilisation subsample with no young children in Tables C.10 and C.11. Parental controls include: individual FEs and parental age FEs. Family controls include: age of youngest child FEs and number of children FEs. Sterilisation restriction + Ex. young restriction: only including families where at least one parent has consulted a doctor regarding sterilisation prior to the announcement of the policy, and excluding families with young children (below the age of 3 throughout the sample period). The sets of controls are discussed in Section 6.3. Standard errors in brackets, clustered at the individual level, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

intensive margins. This motivates two further questions: (1) Who are the women that respond to the reform? And, (2) does the increase in labour supply generally lead to higher levels of income?

9.1. Heterogeneity

First, we look further into the heterogeneity in responses to the reform. For this exercise, we focus on the binary treatment and the number of hours worked, including zeros. We also focus on our best estimates, namely those from model (C) where we control for family characteristics and limit our sample to families where at least one parent has consulted a doctor regarding sterilisation prior to the announcement of the cap on child benefits. We examine heterogeneity in the treatment effect across a range characteristics, namely benefits lost, parental age, age of youngest child, years of education, individual income, and family income. To examine heterogeneity further, we split the treatment group in two along the mean of the characteristic in question, and construct two treatment indicators. The first treatment indicator refers to the treated for whom their value of the characteristic in question falls below the mean, and the second indicator to the treated with values equal to or above the mean. The estimated effects of this exercise are reported in Table 5. For example, in Column 1, the coefficient in the first row refers to treated women who experienced a decrease in child benefits less than 2132 DKK. The coefficient in the second row is the treatment effect for women who experienced a decline in child benefits equal to 2132 DKK or more.

Table 5 shows interesting patterns of heterogeneity in the treatment effect. Firstly, as one would expect, Column 1 reveals that the women experiencing a relatively larger decline in child benefits respond more to the reform. In Column 2, we see that men's labour supply generally does not change in response to the reform. From Columns 3 and 5, we also see that the response to the reform is strongest amongst younger women and those with younger children.

Table 5, Columns 7, 9, and 11 show that the strongest response is amongst women with higher education and relatively lower levels of income. A way to rationalise these findings is that we observe the strongest response amongst women with a relatively high earnings

potential. Interestingly, in Column 10, we also find that men with relatively low earnings respond to the reform, although this effect is relatively imprecisely estimated and only significant at the 5%-level.

9.2. Income and earnings

Because we observe labour supply responses to a reduction in child benefits for mothers, a natural extension is to consider whether the earnings resulting from the increase in labour supply more than offset the decline in income from lost child benefits. Therefore, we consider the reform effect on a range of measures of income in this section. Additionally, in this exercise, we are able to include self-employed in our analyses as their annual income is observed.

Since we now consider annual measures of earnings, we modify our empirical specification as follows:

$$Y_{it} = \beta_0 + \sum_{j=0}^T \beta_1^j 1[j = t] + \beta_2 W_{it} * Intro_t + \beta_4 X_{it} + \alpha_i + \epsilon_{it} \quad (5)$$

t now indicates time in years, and the set of time-varying family controls, X_{it} , now varies at the year-level.¹³ Keep in mind that $Intro_t$ is an indicator equal to one in all periods after the policy introduction. Our main parameter of interest is therefore again β_2 .

We consider the inverse hyperbolic transformation of the various income measures, so our estimated coefficients can be interpreted as percentage changes in income in response to the reform.¹⁴ In contrast

¹³ Included in X_{it} is also a set of year \times 1-digit industry dummies for the self-employed; the reference group are the employees, i.e. those that are not self-employed in a given year. A significant share of the self-employed are farmers, and global food prices fluctuated substantially over the sample period, see e.g. Bellemare (2015). Fluctuations in global food prices translates directly to volatility in profits of self-employed farmers. In addition, other groups of self-employed exposed to variations in earnings due to the Great Recession. The year \times 1-digit industry dummies controls for this exogenous change in earnings of the self-employed. When considering family earnings, a set of year \times 1-digit industry dummies are included for each partner.

¹⁴ All measures of income are inflation-adjusted to 2014-levels.

Table 5

Difference-in-differences: Hours of work per month, by various groups.

Heterogeneity by:	Dependent variable: Hours worked per month											
	Benefits lost		Age		Age of youngest child		Years of education		Ind. income, net tax		Fam. income, net tax	
	Women	Men	Women	Men	Women	Men	Women	Men	Women	Men	Women	Men
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
1.Intro \times 1.Treated < Mean	1.900 [1.250]	0.239 [1.108]	3.788** [1.322]	−1.114 [1.110]	3.545** [1.119]	−0.743 [0.861]	1.805 [1.271]	0.0382 [1.084]	8.085*** [1.295]	1.925* [0.949]	5.135*** [1.269]	1.186 [0.977]
1.Intro \times 1.Treated \geq Mean	2.976** [1.102]	0.0859 [0.872]	1.639 [1.076]	1.093 [0.864]	1.489 [1.273]	1.237 [1.123]	3.576*** [0.994]	0.236 [0.865]	−3.267** [1.005]	−2.932*** [0.879]	−0.142 [1.061]	−1.312 [0.906]
R^2	0.663	0.514	0.663	0.514	0.663	0.514	0.663	0.514	0.663	0.515	0.663	0.514
N	3 205 408	2 575 379	3 205 408	2 575 379	3 205 408	2 575 379	3 205 408	2 575 379	3 205 408	2 575 379	3 205 408	2 575 379
Mean (sample-split condition)	2132	2149	38	41	5	5	15	14	206 665	296 231	485 070	517 109
Female	1	0	1	0	1	0	1	0	1	0	1	0
Parental controls	1	1	1	1	1	1	1	1	1	1	1	1
Family controls	1	1	1	1	1	1	1	1	1	1	1	1
Sterilisation	1	1	1	1	1	1	1	1	1	1	1	1
Ex. young	1	1	1	1	1	1	1	1	1	1	1	1

Notes: Dependent variable: Hours worked per month. This table reports the estimated coefficients from Eq. (3) in Panel A, but we split the treatment group in two along the mean of the characteristic in question based on the value for each individual in December 2010; the means at which the treatment group is divided are reported in the row “Mean (sample-split condition)”. Estimated on monthly observations from October 2008 to December 2014. Monthly hours worked are available from the BFL-register. See details on the full estimation sample in Tables C.8 and C.9, and details on the sterilisation subsample with no young children in Tables C.10 and C.11. Measures of income are available from the IND-register. Income, net of tax also ex. child benefits and refers to “PERINDKIALT_13” - “SKATMVIALT_13” - “KORYDIAL”. Parental controls include: individual FEs and parental age FEs. Family controls include: age of youngest child FEs and number of children FEs. Sterilisation restriction + Ex. young restriction: only including families where at least one parent has consulted a doctor regarding sterilisation prior to the announcement of the policy, and excluding families with young children (below the age of 3 throughout the sample period). The sets of controls are discussed in Section 6.3. Standard errors in brackets, clustered at the individual level, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

to a logarithmic transformation, this allows us to include individuals with zero income. To exclude extreme outliers, we also implement a 98% winsorisation of all income measures.¹⁵

We consider the reform effect on both individual and family income. First, in Table 6, we report the treatment effect on various measures of individual income. In Panel A, we consider aggregate income including public transfers, but excluding child benefits. We find that the income of treated women increases significantly in response to the reform. We also see a positive, but statistically insignificant, increase in income of men. In Panel B, we instead consider earnings, that is, total income excluding public transfers. We find similar positive coefficients, but they are imprecise and insignificantly different from zero. In Panel C, we consider income net of tax, including child benefits. Again, the effects are not statistically significant and the magnitudes are smaller.

However, when we instead consider measures of family income in Table 7, we find similar results in Panels A and B, families increase their earnings in response to the reform, and when considering family income, the effect is very precisely estimated. In Panel C, Column 3, we find that the increase in income carries through when considering income net of tax and including child benefits. The findings from Panels C, Tables 6 and 7, suggest that families tend to increase their net-of-tax income in response to the reform, but whether that happens through an increase in the mother's or father's income varies between families — this would yield imprecise estimates at the individual level, but precise and positive estimates at the family level. Although our previous set of analyses shows that fathers do not significantly adjust their labour supply, the insignificant increase in hours worked after the reform will translate into higher earnings. In addition, fathers can increase their earnings along other margins, e.g. by finding a job with a higher hourly wage.¹⁶

¹⁵ We set values below the 1st percentile equal to the 1st percentile, and similarly, we set values exceeding the 99th percentile equal to the 99th percentile.

¹⁶ All else fixed, if possible, fathers would also take up jobs with a higher wages in absence of the cap on child benefits. However, it may be that they accept a job with a higher wage, but with a lower level of non-pecuniary benefits in response to the reform.

10. Conclusions

A temporary cap on child benefit payments in 2011 led to a non-negligible reduction in child benefits for 3-child families with young children as well as for larger families in Denmark. The differential impact of this policy shift represents an opportunity to assess the causal impact of child benefit programmes on the labour supply of mothers and fathers. Compared to in-work child benefits (e.g. the EITC in US), the economic analysis and interpretation of the cap on child benefits in Denmark is more straightforward as there is no substitution effect between labour supply and child benefits; the levels of child benefits depend on family composition, not income. Therefore, we can isolate the income effect of the reform. We find that a reduction in child benefits increases the labour supply of mothers, but not of fathers. Even after controlling for fertility-related family characteristics, we find a significant increase in mothers' labour supply after the introduction of the reform. This is in line with the existing literature, which highlights that women's labour supply is relatively elastic when compared to men, e.g. Evers et al. (2008).

We find that mothers in affected families respond to the reform at the intensive margin, but that the strongest response to the reform is at the extensive margin. We confirm our results by using data on parents' sterilisation consultations with doctors, a common procedure in Denmark. The advantage of the data on sterilisation is that we can limit our analyses to families who had finalised their fertility decisions prior to the announcement of the reform. By limiting our sample to families in which at least one parent has consulted a doctor on sterilisation, we are able to hold preferences for additional children fixed across the treatment and control groups, which otherwise are likely to differ between families with different numbers of children. If, due to a higher number of children, our treatment group has stronger preferences for leisure/time with children than our control group conditional on fertility-related characteristics, our labour supply responses should provide a lower bound for the group of families with fewer children.

Using our group of families that have observably finalised fertility, i.e. consulted a doctor regarding sterilisation, we estimate labour supply responses to the reform conditional on similar preferences for additional children between our treatment and control groups. When considering the total labour supply response including responses at

Table 6
Difference-in-differences, binary treatment: Individual income.

	Women			Men		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Total income, ex. child benefits						
1.Intro × 1.Treat	0.00190 [0.00315]	0.00921** [0.00343]	0.0520*** [0.0152]	0.0127** [0.00431]	0.00775 [0.00466]	0.0348 [0.0193]
R ²	0.716	0.717	0.728	0.662	0.662	0.661
Panel B: Total income, ex. public transfers and child benefits						
1.Intro × 1.Treat	0.0560*** [0.0132]	0.0221 [0.0143]	0.0367 [0.0540]	0.0173 [0.0115]	0.00718 [0.0125]	0.0862 [0.0512]
R ²	0.683	0.686	0.741	0.607	0.607	0.621
Panel C: Total income, net of tax, inc. child benefits						
1.Intro × 1.Treat	0.00556*** [0.00147]	−0.0157*** [0.00158]	−0.00308 [0.00620]	0.00756*** [0.00215]	−0.0123*** [0.00230]	0.00125 [0.00947]
R ²	0.762	0.764	0.774	0.701	0.703	0.702
N	1 819 230	1 819 230	315 474	1 569 337	1 569 337	268 579
Female	1	1	1	0	0	0
Parental controls	1	1	1	1	1	1
Family controls	0	1	1	0	1	1
Sterilisation	0	0	1	0	0	1
Ex. young	0	0	1	0	0	1

Notes: Dependent variable: various measures of income. This table reports the estimated coefficients from Eq. (5). Estimated on yearly observations from 2008 to 2014. Measures of income are available from the IND-register. All measures of income are inflation-adjusted to 2014-levels. Total income, ex. child benefits, refers to “PERINDKIALT_13” - “KORYDIAL”. Total income, ex. public transfers and child benefits, refers to “PERINDKIALT_13” - “OFF_OVERFORSEL_13” (OFF_OVERFORSEL_13 includes “KORYDIAL”). Total income, net of tax, inc. child benefits, refers to “PERINDKIALT_13” - “SKATMVIALT_13” - “KORYDIAL” +0.5× our precise measures of yearly child benefits (“KORYDIAL” provides an imprecise weighted average over two years). Child benefits are family-specific, so we assign 50% to each parent. See details on the full estimation sample in Tables C.8 and C.9, and details on the sterilisation subsample with no young children in Tables C.10 and C.11. Parental controls include: individual FEs, parental age FEs, and a set of year × 1-digit industry dummies for the self-employed. Family controls include: age of youngest child FEs and number of children FEs. Sterilisation restriction + Ex. young restriction: only including families where at least one parent has consulted a doctor regarding sterilisation prior to the announcement of the policy, and excluding families with young children (below the age of 3 throughout the sample period). The sets of controls are discussed in Section 6.3. Results for the same outcomes, but with a continuous treatment variable, are available in Table C.18. Standard errors in brackets, clustered at the individual level, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

both the extensive and intensive margins, our results suggest an upper bound of the female labour supply elasticity with respect to net income of -2.700 ,¹⁷ and a lower bound of -0.900 .¹⁸ If we instead consider only the intensive margin, which is often done in the literature when estimating income rather than wage elasticities, the estimated elasticities should be adjusted by a factor of 0.434, giving a range between -1.161 and -0.387 .¹⁹ The upper bound estimates assume that families had anticipated that the cap on child benefits would only be in place for a single year, and the lower bound estimate takes into account that families expected an increased income loss in the years to come due to the gradual implementation of the policy. Although the policy was ultimately only in place in 2011, families were unable to predict this repeal of the policy. Therefore, we believe that our lower bound

estimates reflect our best estimates of female labour supply elasticities with respect to net income.

Even our lower bound estimates of female labour supply elasticities are large in absolute values compared to the majority of estimates in the existing literature, which often find values near zero for married and cohabiting mothers (see e.g. Hotz and Scholz, 2003; McClelland and Mok, 2012, for reviews). On the other hand, Eissa and Hoynes (2004, 2006b) find statistically significant and negative female labour supply elasticities, as they document a substantial reduction in maternal labour supply with the expansion of the EITC for married couples in the US. Eissa and Hoynes (2006b) estimate an income elasticity between -0.04 and -0.36 at the intensive margin, which is close to our best estimate of -0.387 for that margin. In addition, our estimated labour supply elasticities are close to those estimated for single mothers and low-income mothers in the literature.

There are a number of reasons why we may expect larger female labour supply elasticities in absolute values in our setting compared to those in the existing literature. First, we consider families who have finalised fertility. If mothers do not expect any additional children, they can assume that time spent on child care will only decrease over time, and therefore, they may be more inclined to increase their labour supply in response to an income loss. Second, the 2011 cap on child benefits considered in this paper is the first substantial deterioration in the provision of child benefits in Denmark; a very generous benefit system that had gradually been expanded since the 1980s (see Appendix A for details). Thus, if families are risk averse, they may be less likely to base their labour supply decisions on child benefits income after the reform if future levels of benefits are perceived as more uncertain.

¹⁷ We use estimates from model (C), see Section 6.3 for details, to estimate income elasticities. Using data from Tables 2 (Column 3) and C.10, we have that $\Delta h/\Delta I = (2.688/119.368)/(-2103.134/(210765.315 + 41309.403))$. This assumes that child benefits are perceived as a loss in mothers' income. Since child benefits at the time of the reform were paid to mothers' bank accounts by default, we believe this is appropriate.

¹⁸ We use estimates from model (C), see Section 6.3 for details, to estimate income elasticities. Using data from Tables 2 (Column 3) and C.10, we have that $\Delta h/\Delta I = (2.688/119.368)/((3 \times -2103.134)/(210765.315 + 41309.403))$. This lower bound considers that most families expected a future income loss from the cap on child benefits equal to three times the loss of benefits in 2011. Please see the discussion of the gradual implementation of the policy in Appendix A and an illustration of this in Fig. 2.

¹⁹ The intensive margin response account for $0.434 = 1.167/2.688$ of the total labour supply response, see Tables 2 and 4, Panels A, Column 3.

Table 7
Difference-in-differences, binary treatment: Family income.

	Women			Men		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Total income, ex. child benefits						
1.Intro \times 1.Treat	0.00960*** [0.00145]	0.00944*** [0.00161]	0.0417*** [0.00641]	0.00895*** [0.00155]	0.0119*** [0.00167]	0.0446*** [0.00670]
R^2	0.822	0.823	0.818	0.805	0.806	0.795
Panel B: Total income, ex. public transfers and child benefits						
1.Intro \times 1.Treat	0.0346*** [0.00379]	0.0243*** [0.00415]	0.0700*** [0.0171]	0.0305*** [0.00363]	0.0215*** [0.00393]	0.0580*** [0.0164]
R^2	0.711	0.713	0.724	0.689	0.691	0.688
Panel C: Total income, net of tax, inc. child benefits						
1.Intro \times 1.Treat	0.00969*** [0.00112]	-0.00652*** [0.00122]	0.0182*** [0.00492]	0.00586*** [0.00120]	-0.00635*** [0.00128]	0.0183*** [0.00514]
R^2	0.794	0.796	0.789	0.771	0.773	0.761
N	1 819 230	1 819 230	315 474	1 569 337	1 569 337	268 579
Female	1	1	1	0	0	0
Parental controls	1	1	1	1	1	1
Family controls	0	1	1	0	1	1
Sterilisation	0	0	1	0	0	1
Ex. young	0	0	1	0	0	1

Notes: Dependent variable: various measures of income. This table reports the estimated coefficients from Eq. (5). Estimated on yearly observations from 2008 to 2014. Measures of income are available from the IND-register, and they are aggregated at the family-level. All measures of income are inflation-adjusted to 2014-levels. Total income, ex. child benefits, refers to “PERINDKIALT_13” - “KORYDIAL”. Total income, ex. public transfers and child benefits, refers to “PERINDKIALT_13” - “OFF_OVERFORSEL_13” (OFF_OVERFORSEL_13 includes “KORYDIAL”). Total income, net of tax, inc. child benefits, refers to “PERINDKIALT_13” - “SKATMVIALT_13” - “KORYDIAL” + our precise measures of yearly child benefits (“KORYDIAL” provides an imprecise weighted average over two years). See details on the full estimation sample in Tables C.8 and C.9, and details on the sterilisation subsample with no young children in Tables C.10 and C.11. Parental controls include: individual FEs, parental age FEs, and a set of year \times 1-digit industry dummies for the self-employed. Family controls include: age of youngest child FEs and number of children FEs. Sterilisation restriction + Ex. young restriction: only including families where at least one parent has consulted a doctor regarding sterilisation prior to the announcement of the policy, and excluding families with young children (below the age of 3 throughout the sample period). The sets of controls are discussed in Section 6.3. Results for the same outcomes, but with a continuous treatment variable, are available in Table C.19. Standard errors in brackets, clustered at the individual level, $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

A perceived change in the stability of the child benefit system may induce a “cultural shift” amongst the affected families; that is, a shift away from the culture in which child benefits were perceived as a reliable source of income. In response to this shift, affected families may increase their labour supply to limit the impact of potential future child benefit income losses. This also points to a broader question around the interpretation of parameter estimates driven by policy changes that may foster long-term cultural shifts. In the case of the policy change we consider, the fact that it was enacted soon after the Great Recession could potentially solidify a cultural shift in response to the reform. Third, most of the existing studies that evaluate labour supply effects of unconditional child benefits exploit policies that increase the level of child benefits. However, income effects are not necessarily symmetric; families may react differently to an income gain and an income loss of similar magnitudes, for example due to credit constraints (see e.g. Christelis et al., 2019, for a discussion of asymmetric responses to positive and negative income shocks). If a family is credit constrained, they would most likely adjust their labour supply when facing an income loss, but not necessarily when experiencing an income gain. Fourth, we should also note that we estimate the effect of the reform in Denmark, a country where the labour force participation of women – including married mothers – is very high relative to other countries (OECD, 2020). In other countries and cultures, where the institutions and norms that affect maternal supply differ, the effects of similar reforms would most likely differ from our estimates.

We also consider heterogeneity in the labour supply response within our treatment group. As expected, we find that mothers in families that face a larger reduction in child benefits also respond more to the

reform. In addition, we find that young mothers with a high level of education and a relatively low level of income respond relatively more to the reform. We rationalise these findings as mothers with the highest earnings potential being more likely to increase their labour supply in response to the reform.

Furthermore, the effects on mothers’ labour supply are persistent, even after the repeal of the reform. There are at least two explanations of the persistence of the effects: (1) The costs of entry to the labour market/the costs of increasing work hours have already been borne by the mothers, e.g. by enrolling children in daycare and kindergarten. (2) The reform introduced uncertainty about future child benefit payments, even after its repeal, as the generosity of the payments received increased political and public awareness. As discussed above, a perceived change in the stability of the child benefit system could trigger a “cultural shift” amongst the affected families with long-term effects. We also find that the income of affected families increased in response to the reform because the reduction in child benefits was ultimately temporary (due to the repeal of the reform from 2012), while the labour supply effects were sustained. However, the increase in income appears to be driven by both mothers and fathers in affected families, suggesting that fathers respond on other margins than labour supply, e.g. by changing jobs.

In terms of policy implications, our results complement existing evidence on the EITC from the US, e.g. Meyer and Rosenbaum (2000), which generally finds positive employment effects of in-work benefits on single mothers. The universal child benefit system in Denmark appears to have the opposite effect for women in two-parent families. Thus, though the programme in Denmark is quite different, our results

support the conclusion of Eissa and Hoynes (2004, p. 1931) who found that in the US, “the EITC is effectively subsidizing married mothers to stay home...” Depending on policymakers’ objectives, alternative policies could be developed, e.g. to target the child pay penalty or part-time pay penalty of mothers. However, other outcomes, such as the wellbeing and poverty of children and mothers, should, of course, also be considered in the design of such policies.

Declaration of competing interest

I can confirm that the authors, Mathias Fjællegaard Jensen and Jack Blundell, have no affiliation with or involvement in any organization or entity with any financial interest, or non-financial interest in the subject matter or materials discussed in this manuscript.

Data availability

The authors do not have permission to share data.

Appendices A to C. Policy background, conceptual framework, and additional results

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jpubeco.2023.105049>.

References

- Almlund, M., 2018. Fertility and child benefits. In: VIVE - Danish National Centre for Social Research. mimeo.
- Angrist, J.D., Pischke, J.-S., 2008. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press, Princeton, N.J..
- Bellemare, M.F., 2015. Rising food prices, food price volatility, and social unrest. *Am. J. Agric. Econ.* 97 (1), 1–21.
- Blau, F.D., Kahn, L.M., 2017. The gender wage gap: Extent, trends, and explanations. *J. Econ. Lit.* 55 (3), 789–865.
- Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L., Van Rooij, M., 2019. Asymmetric consumption effects of transitory income shocks. *Econom. J.* 129 (622), 2322–2341.
- Cogan, J.F., 1981. Fixed costs and labor supply. *Econometrica* 49 (4), 945–963.
- Corinth, K., Meyer, B., Stadnicki, M., Wu, D., 2021. The anti-poverty, targeting, and labor supply effects of replacing a child tax credit with a child allowance. In: NBER Working Paper, No. 29366.
- Dahl, G.B., Lochner, L., 2012. The impact of family income on child achievement: Evidence from the earned income tax credit. *Amer. Econ. Rev.* 102 (5), 1927–1956.
- Eissa, N., Hoynes, H.W., 2004. Taxes and the labor market participation of married couples: the earned income tax credit. *J. Public Econ.* 88 (9), 1931–1958.
- Eissa, N., Hoynes, H.W., 2006a. Behavioral responses to taxes: Lessons from the EITC and labor supply. In: *Tax Policy and the Economy*, Volume 20. National Bureau of Economic Research, Inc, pp. 73–110.
- Eissa, N., Hoynes, H.W., 2006b. The hours of work response of married couples: Taxes and the earned income tax credit. In: Agell, J., Sørensen, P.B. (Eds.), *Tax Policy and Labor Market Performance*. The MIT Press.
- Eissa, N., Liebman, J.B., 1996. Labor supply response to the earned income tax credit. *Q. J. Econ.* 111 (2), 605–637.
- Evers, M., Mooij, R., Vuuren, D., 2008. The wage elasticity of labour supply: A synthesis of empirical estimates. *Economist* 156 (1), 25–43.
- Fernández, R., Fogli, A., 2009. Culture: An empirical investigation of beliefs, work, and fertility. *Am. Econ. J.: Macroecon.* 1 (1), 146–177.
- Financial Times, 2021. US embarks on first national child allowance experiment.
- Francesconi, M., van der Klaauw, W., 2007. The socioeconomic consequences of “In-Work” benefit reform for British lone mothers. *J. Hum. Resour.* 42 (1).
- Gelber, A.M., Mitchell, J.W., 2012. Taxes and time allocation: Evidence from single women and men. *Rev. Econom. Stud.* 79 (3), 863–897.
- González, L., 2013. The effect of a universal child benefit on conceptions, abortions, and early maternal labor supply. *Am. Econ. J. Econ. Policy* 5 (3), 160–188.
- Hener, T., 2016. Unconditional child benefits, mothers’ labor supply, and family well-being: Evidence from a policy reform. *CESifo Econ. Stud.* 62 (4), 624–649.
- Hotz, V.J., Scholz, J.K., 2003. The earned income tax credit. In: Moffitt, R.A. (Ed.), *Means-Tested Transfer Programs in the United States*, No. January. University of Chicago Press, pp. 141–197.
- Hotz, V.J., Scholz, J.K., 2006. Examining the effect of the earned income tax credit on the labor market participation of families on welfare. In: NBER Working Paper, No. 11968. National Bureau of Economic Research.
- Kleven, H., 2019. The EITC and the extensive margin: A reappraisal. In: NBER Working Papers, No. 26405. National Bureau of Economic Research, Inc.
- Kleven, H., Landais, C., Søgaard, J.E., 2019. Children and gender inequality: Evidence from Denmark. *Am. Econ. J. Appl. Econ.* 11 (4), 181–209.
- Kuka, E., Shenhav, N., 2020. Long-run effects of incentivizing work after childbirth. In: NBER Working Paper, Vol. 27444.
- McClelland, R., Mok, S., 2012. A review of recent research on labor supply elasticities. In: Working Paper.
- Meyer, B.D., 2010. The effects of the earned income tax credit and recent reforms. *Tax Policy Econ.* 24 (1), 153–180.
- Meyer, B.D., Rosenbaum, D.T., 2000. Making single mothers work: Recent tax and welfare policy and its effects. *National Tax J.* 53 (4 PART 2), 1027–1059.
- OECD, 2020. OECD Family Database: LMF1.2. Maternal employment rates Definitions. Technical Report November, pp. 1–8.
- Raschke, C., 2016. The impact of the german child benefit on household expenditures and consumption. *Ger. Econ. Rev.* 17 (4), 438–477.
- Ray, R., Gornick, J.C., Schmitt, J., 2010. Who cares? assessing generosity and gender equality in parental leave policy designs in 21 countries. *J. Eur. Soc. Policy* 20 (3), 196–216.
- Schanzenbach, D.W., Strain, M.R., 2021. Employment effects of the earned income tax credit: Taking the long view. *Tax Policy Econ.* 35 (1), 87–129.
- Tamm, M., 2010. Child benefit reform and labor market participation. *Jahrb. Nationalökonomie Stat.* 230 (3), 313–327.